Latent spatial and temporal structures in fluvial river temperatures sensed using a novel wireless sensor network

By

Scott G Burman
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Ecology

in the

COLLEGE OF AGRICULTURAL AND ENVIRONMENTAL SCIENCES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

John Largier, Chair

Dipak Ghosal

Nann Fangue

Committee in Charge

2023

Table of contents

Abstract

In an effort to describe water temperatures at scales relevant to organisms living in the lower Yuba River (LYR) in California's Central Valley, we designed and deployed a wireless sensor network in November 2018 through May 2019. Temperatures were measured along a 3 km study reach, across the channel, and into off-channel areas. To capture diel and seasonal fluctuations, sensors were sampled quarter-hourly for six months.

Chapter 1 describes the wireless sensor design, field deployment, data cleaning, and raw data. The deployment adopted event-based software on MSP430 micro-controllers with 433 MHz radio and minimized the networking duty-cycle. To address link failures, we included network storage. As the network lacked real-time clocks, data were timestamped at the destination, which – with the network design – yielded timestamp inaccuracies that were re-aligned algorithmically. We collected over six months of temperature data from 35 sensors across seven nodes. Of the packets collected, we identified 21% as being incorrectly time-stamped and were able to re-align 41% of these incorrectly timestamped packets.

In chapter 2, we use these data to consider the dominant trends in fluvial temporal dynamics. Generally, river waters warm as they move downstream. We had three research questions: At what temporal scales do fluvial temperatures vary most? Do hyporheically connected side-channels experience warming, or do connections to the mainstem limit these effects? Over short reaches of relatively cold water, is the downstream warming trend still prevalent? We created several Bayesian models, which segmented the river geographically, but were not spatially explicit. We found that models that incorporated the interactive effect of a diel and day-of-sampling effects with random walk priors yielded the best fit. Within the study reach, model results indicated that temporal effects dominated the temperature variance, with diel-effects less important than multi-day effects. We found that even with strong hyporheic connections, the sampled side channel was warmer than the main-channel. We also found only minimal downstream warming within our 3 km study reach.

In chapter 3, we describe the spatial structures within a side channel of the study reach. The primary research question addresses whether the spatial structures of river temperatures (longitudinal and lateral) are static or if they change over time, either by time of day or by day of year. We then question if the time-dependence of these spatial structures is related to river flow rate and/or surface heating, indexed by daily solar radiation. The lateral and longitudinal distances between sensors were used to generate semi-variograms, calculated for specific time-of-day and specific day-of-year time bins. We found that the spatial structures of fluvial temperatures were time varying. Diel changes in lateral structure in temperature corresponded with the diel cycle in solar radiation, while day-to-day and seasonal changes in lateral structure corresponded with changes in river flow rate. Longitudinal structure did not exhibit a coherent diel cycle, but we did observe seasonal changes in semi-variogram slope with marked longitudinal structure during low warm-water flows in November, high cold-water flows in February, and with the return of strong diurnal warming in May.

TempMesh – A Flexible Wireless Sensor Network for Monitoring River Temperatures

SCOTT G. BURMAN, JINGYA GAO, GREGORY B. PASTERNACK, and
NANN A. FANGUE, University of California, Davis
PAUL CADRETT, U.S. Fish and Wildlife Service
ELIZABETH CAMPBELL, U.S. Army Corps of Engineers
DIPAK GHOSAL, University of California, Davis

For a Chinook salmon restoration project in the lower Yuba River in California, we designed and deployed a wireless sensor network to monitor river temperatures at micro-habitat scales. The study required that temperatures be measured along a 3 km study reach, across the channel, and into off-channel areas. To capture diel and seasonal fluctuations, sensors were sampled quarter-hourly for the full duration of the six-month juvenile salmon winter residency. This sampling duration required that nodes minimize power-use. We adopted event-based software on MSP430 micro-controllers with 433 MHz radio and minimized the networking duty-cycle. To address link failures, we included network storage. As the network lacked real-time clocks, data were timestamped at the destination. This, coupled with the storage, yielded timestamp inaccuracies, which we re-aligned using a novel algorithm. We collected over six months of temperature data from 35 sensors across seven nodes. Of the packets collected, we identified 21% as being incorrectly timestamped and were able to re-align 41% of these incorrectly timestamped packets. We collected temperature data through major floods, and the network uploaded data until the flood destroyed the sensors. The network met an important need in ecological sampling with ultra-low power (multi-year battery life) and low-throughput.

CCS Concepts: • **General and reference → Empirical studies**; • **Hardware → Sensor applications and deployments**; • **Computer systems organization → Sensor networks**; • **Networks → Network protocol design**; **Sensor networks;**

Additional Key Words and Phrases: River temperature monitoring, wireless sensor network, fluvial temperatures, network storage, power efficiency, timestamp alignment, sensor deployment, fish, salmon

Scott G. Burman and Jingya Gao contributed equally to this research.

This work was supported by a cooperative agreement (F16AC00735) with the Anadromous Fisheries Restoration Program within the US Fish and Wildlife Service. The USDA National Institute of Food and Agriculture support Greg Pasternack (Hatch: CA-D-LAW-7034-H) and Nann Fangue (Hatch: CA-D-WFB-2098-H).

Authors' addresses: S. G. Burman (corresponding author), J. Gao, G. B. Pasternack, N. A. Fangue, and D. Ghosal, University of California, Davis, 1 Shields Avenue, Davis, CA, 95616; emails: {sgburman, aligao, gpast, nafangue, dghosal}@ucdavis.edu; P. Cadrett, U.S. Fish and Wildlife Service, 850 S Guild Ave Suite 105, Lodi, CA, 95240; email: paul_cadrett@fws.gov; E. Campbell, U.S. Army Corps of Engineers, San Francisco District, 450 Golden Gate Avenue, 4th Floor, San Francisco, CA, 94102; email: elizabeth.a.campbell@usace.army.mil.

1550-4859/2022/12-ART15
https://doi.org/10.1145/3542697

ACM Reference format:
Scott G. Burman, Jingya Gao, Gregory B. Pasternack, Nann A. Fangue, Paul Cadrett, Elizabeth Campbell, and Dipak Ghosal. 2022. TempMesh – A Flexible Wireless Sensor Network for Monitoring River Temperatures. *ACM Trans. Sensor Netw.* 19, 1, Article 15 (December 2022), 28 pages.
https://doi.org/10.1145/3542697

1 INTRODUCTION

Anadromous fish like salmon swim from oceans into rivers where they spawn. While their time in rivers is much shorter than their oceanic life-stages, the in-river stages are extremely energetically costly: Adults swim for miles against currents (without feeding), generate gametes, and then distribute them by spawning. Then, the next generation's juveniles hatch, grow in size by orders of magnitude, avoid predators, find and consume food, and finally smoltify—which is a physiological change from juvenile parr to smolt, to prepare for downstream migration. These anadromous fish must be adapted to the oceans where they live, which are fairly homogeneous in temperature. At the same time, they must retain (and pass on to their progeny) the ability to expend considerable energy and thrive in rivers with unique temperature profiles. To reach our study system, the **lower Yuba River (LYR)**, the fish must first enter the San Francisco Bay, move up the Sacramento River and then into the Feather River, migrate up the Feather River, then move into the Yuba River. This migration is a particular challenge, as the Feather River is considerably warmer than the very cold Yuba River. This study system is therefore ideal for addressing questions about thermal tolerances of salmonids. The diversity of temperatures these fish traverse is extreme, and we sought insights into their adaptation to the different rivers that they traverse to reach spawning sites, and where juveniles grow and smoltify. Beyond these adaptive or evolutionary energetic impacts of temperature, we endeavored to place river temperature into an ecological context by assessing temperatures in micro-habitats. Micro-habitats in this context are at scales below the river-reach, which required sensors dispersed both across the channel and longitudinally along the length of the river.

Fluvial river temperatures are dynamic and complex. Two trends tend to dominate river temperatures: They generally increase moving downstream from headwaters to mouth [1] and are related-to (but not forced by) air-temperatures [2]. Yet, these trends vary both seasonally and daily [3]. Solar warming is a strong predictor of river temperature [4], and as a result, vegetative shading can have intense local effects on micro-habitat temperatures [5]. Finally, flows through sediments, which are shaded from the sun and have no convective warming from air can maintain relatively consistent temperatures over vast distances [6]. Integrating these various inputs makes temperature predictions in rivers challenging. Modeling river temperatures was found to be feasible, but computationally difficult, particularly at fine-scales. Physics-based models were less effective than interpolative models [7, 8].

We found that existing sensing methodologies (FLIR [9] and fiber-optics [10]) could not meet the (spatial and temporal) sampling density and extent that we required to characterize juvenile micro-habitats and instead opted to build a wireless sensor network to collect real-time temperature data.

Wireless sensor networks (WSNs) [11, 12] have been used in monitoring applications in similarly harsh applications, including in the early detection of forest fires [13], in actuating applications such as precision agriculture [14], to measure and control energy usage in smart homes [15, 16], for high-resolution spatio-temporal monitoring of underwater environments [17], and in tracking applications such as animal telemetry [18]. In a typical scenario, sensors are distributed over a geographical area and the data from these sensors are—using a network of relay nodes—aggregated at a gateway node, which in turn uploads the data to a server for processing and analysis.

Multiple environmental factors influenced the design of the sensors and the wireless mesh network that we created. First, the temperature sensors needed to be able to survive changing flows and river conditions. Second, the river on which sensors were installed was difficult to access and there was no access to reliable power. Consequently, sensors and wireless networking nodes were battery-powered. Once deployed, changing batteries at remote node locations was impractical and costly. It was therefore imperative that we optimize the software to minimize power use. There were also challenges to consider when placing nodes: topography and geography limited the places where sensors and nodes could be placed. Finally, hikers, poachers, and animals (beavers and bears) damaged equipment. These constraints limited the density and positioning of relay nodes, which meant that they were spread over larger distances, which rendered link quality more susceptible to failure during inclement weather. To solve these issues, we introduced a network layer function to store data in intermediate nodes until network links were re-established.

This is a **practical implementation paper**, in which we describe the design, implementation, and deployment of a wireless mesh network to gather spatial and temporal river-temperature data from the lower Yuba River at scales appropriate to assess juvenile salmonid micro-habitat. The key contributions of this article include:

(1) Design considerations of the data transfer protocol and network architecture, which include a network storage function robust to intermittent link failure;
(2) Node software architecture based upon event driven proto-threads, designed explicitly to minimize power use;
(3) Experimental results that quantify the node power use and an evaluation of the efficacy of the architecture and implemented low power modes;
(4) A methodology for correcting timestamps of data with an incremental identifier, collected via a network with unknown transmission delays.

The remainder of this article is organized as follows: In Section 3, we discuss some of the technological choices we made and the constraints imposed by these choices. In Section 4, we describe the network architecture and the data transfer protocol. We present how the network storage function was implemented to account for intermittent link failure. In Section 5, we describe the software architecture of the node based on proto-threads and the methods that were implemented to minimize power. In Section 6, we present the field results and discuss the important issue of timestamp alignment. In Section 7, we characterize what we have learned from this work and how it contributes to *in situ* environmental sensing more broadly. In Section 8, we discuss the related work, and in Section 9, we give the conclusions and discuss the future work.

2 BACKGROUND AND MOTIVATION

In 2014, river managers and experts working on the Yuba Accord Management Team for the lower Yuba River in California's Central Valley sought to enhance Chinook salmon populations by targeting the juvenile life-stage for in-river restoration efforts. There was hope that restoration work could modify the river to better meet the energetic needs of juvenile Chinook salmon, which were historically abundant in the Yuba River [19]. Energy balance is dictated by exertion and consumption. To improve the available habitat for the energetics of an organism, the system must either have more abundant or more nutritious food (higher consumption), or require that the organism use less energy in the system (less exertion). Energy use of salmon is largely dependent upon water temperature. To find the optimal temperatures for salmonids, laboratory experiments have considered the relationship between water temperature and aerobic scope (the difference in the fish's energy-use at maximum exertion and at rest). These studies found that salmonids were adapted to the temperatures they experienced in their home rivers [20, 21].

We determined that to meaningfully sample the temperature experienced by juvenile salmonids in the lower Yuba River, we would need hourly temperature data for a 3 km reach of this multi-threaded river. To capture the full juvenile life-stage, this sampling needed to survive in-river for at least six months. We also knew that sensors needed to be able to record and report data even when we could not physically access the river during flood events. Monitoring needed to be highly resolved, both in space and time. Specifically, it was important that the temperature be monitored both longitudinally (along the river, with sensors spaced hundreds of meters apart) and across the channel (with sensors spaced a few meters apart). We opted to sample every 15 minutes to provide multiple samples per hour.

3 TECHNOLOGY

3.1 Temperature Sensing

There exist a few sensing technologies capable of the type of data collection required for the project. **Forward-looking infrared (FLIR)** cameras can be used to image temperatures at the necessary spatial density, either aerially or from channel margins. Unfortunately, achieving the required spatial extent was infeasible, as it would require that we install dozens of FLIR cameras or move them constantly. They also can only image the water's surface [9]. FLIR deployments also fail to meet the temporal extent needed for this study, as they would require considerable maintenance and power over the course of the six-month juvenile residency period.

Fiber-optic networks can sample large distances (~13 km) at a time, with about 1 meter spatial resolution [10]. As a result, fiber optic temperature sensing met the necessary spatial density and extent for our work, with the fiber optics repeatedly crossing the river in a zig-zag pattern. Unfortunately, fiber-optic temperature sensing was not feasible for this study due to the required lasers, which use considerable amounts of power and the need for continuous, on-going upkeep and maintenance. This was impractical given the access issues of the study sites and the sampling duration.

It became clear that the best option for temperature monitoring was the use of electronic temperature sensors. We initially considered data-logging sensors, but moved away from these because the lower Yuba River experiences extreme flood events that can transport trees, large boulders, and other detritus down-river. These events can damage, bury, or dislodge temperature sensors. Since there is no way to remotely recover data from low-cost *in situ* logging sensors, destructive events eliminate all data since the previous download, and data from flood events is lost.

Ultimately, we opted to construct a wireless sensor network that sensed data using off-the-shelf temperature sensors. We selected the Campbell Scientific CS225-L SDI-12 temperature sensor cables, which were environmentally sealed and protected all sensor hardware by housing it in a thick, sealed cable assembly. These assemblies included individually addressable, digital temperature sensors, allowing for long cable runs without concern about voltage loss or other drift that could alter sensor readings. These sensors were sufficiently accurate for the study ($\pm 0.2^oC$). Sensor assemblies were designed to contain five sensors, evenly spaced either 5 m or 10 m apart along the length of the cable, depending upon the river width at a particular location. Sensor cables also had between 10 m and 100 m of additional cabling on the end to allow for transmitters to be concealed in riparian foliage. These sensors were tethered to anchors buried in the riverbed. Sensors were initially designed to withstand large changes in the flow of the river, with sacrificial anchors, which would detach under destructive flows and allow the sensor cables to survive. Despite these precautions, it proved nearly impossible to design the sensor systems to survive impacts from massive detritus such as logs carried down by the river during the highest of flows.

3.2 Wireless Mesh Network

The Campbell Scientific CS225-L temperature sensors were integrated with a wireless mesh network built using inexpensive, low-power devices. Each device was equipped with a radio with reasonable range to aggregate the temperature data from the sensor strings to one aggregation point from which the data could be back-hauled to a server on the University of California, Davis, campus using a cellular connection.

There are four main technologies for implementing wireless mesh networks [22], based on (1) IEEE 802.11s, (2) **Bluetooth Low Energy (BLE)**, (3) IEEE 802.15.4, and (4) LoRa (and the corresponding MAC layer protocol, LoRaWAN). The key advantage of devices built using these technologies is that they come with protocols for node discovery, topology formation, and data forwarding and routing. These technologies also have high bandwidth. The range for BLE-based devices is about 30 m. While BLE's relatively high bandwidth at this range is optimal for its intended application, it was infeasible for the present study, as 30 m node ranging was inadequate and would have required far more mesh nodes than could be installed in the given terrain. IEEE 802.11s [23], IEEE 802.15.4 [24], and LoRa [25] are designed for **Internet of Things (IOT)** applications in industrial setting or Smart City applications where either power is readily available or batteries can be easily recharged or replaced. These technologies were not designed to minimize power use. We wanted to use 433 MHz band because of its superior propagation characteristics outdoors compared to 2.4 GHz or 900 MHz, despite the negative effects of having a larger Fresnel zone. As a result, the 433 MHz band has gained momentum for machine-to-machine communications using low-power wireless technologies [26]. DASH7 and IEEE 802.15.4f are two standards that have been developed for the 433 MHz band. At the time of the deployment, LoRa had not yet added support for the 433 MHz band. Our application required reliable data transmission in an energy-deficient environment. We prioritized power savings and data fidelity over bandwidth, latency or other considerations, because the raw temperature data occupied only a few bytes, which meant each sensor node could generate small packets every 15 minutes.

We opted to use wireless communication operating at 433 MHz, because the frequency struck a reasonable balance between range and penetration. It also was an unregulated frequency band—requiring only a **Federal Communications Commission (FCC)** experimental license. Using low-power transmitters like ours, operating at 433 MHz, 500 m range was achievable with simple, low-cost whip antennae. At this range, 433 MHz transmissions can penetrate vegetation while reflecting off of solid surfaces. In a river, the steep embankments leading up to the floodplains and levees can create something of an echo chamber, which we found to enhance transmission ranges.

3.3 Hardware

The nodes of the wireless mesh network were implemented using WizziMotes [27], an IoT device built on top of the **Texas Instruments (TI)** CC4305137, an MSP430-based **micro-controller unit (MCU)** with integrated 433 MHz wireless networking. The MSP430 family of MCUs are ideal for this application, as they are low-power and can be switched to ultra-low power modes, which significantly extend battery life [28, 29]. We used WizziMote because they were among the only devices available at the time that operated in the 433 MHz band and implemented some aspects of the DASH7 protocol [27, 30].

As a **System-on-Chip (SOC)**, MSP430 provides integrated peripherals for a variety of battery-operated wireless applications. The operating modes take into account three different needs: ultra-low power, speed and data throughput, and minimization of individual peripheral power consumption [28, 29]. Built with **Low Power Modes** (**LPMs**; LPM0: 80μA, LPM2: 6.5μA, LPM3: 2μA, LPM4: 1μA), MSP430 can preserve energy by shutting down respective clocks on the processing chip and

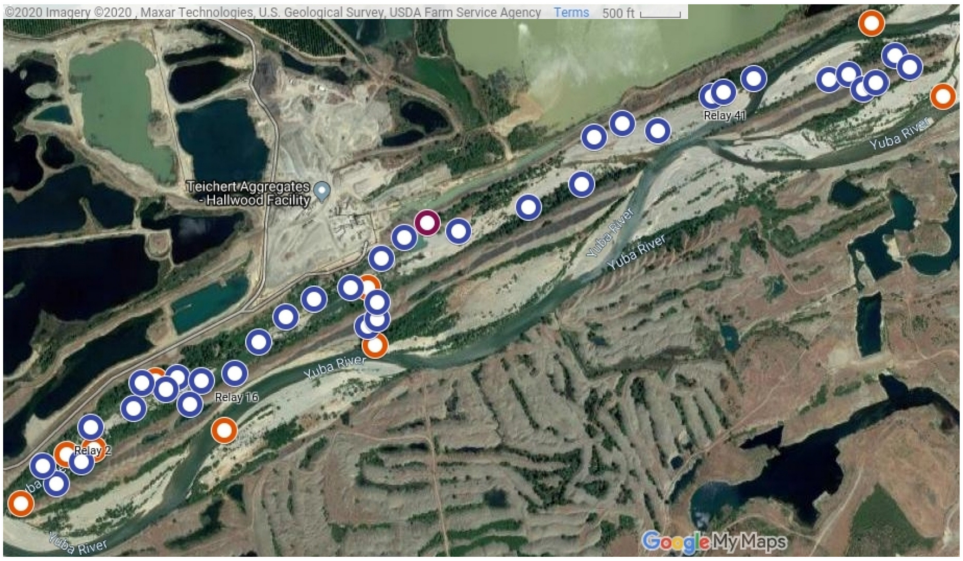

Fig. 1. Map of the network deployed in November 2018, which remained functional through May 2019 on the 3 km study reach of the lower Yuba River, immediately downstream from Daguerre Point Dam (not featured: just upstream of top-right of map) near Marysville, in California's Central Valley. The orange nodes are the sensor nodes, the blue nodes are the relay nodes, and the magenta node is the gateway node. Map generated using Google My Maps, and is 3.7 km by 2.1 km.

later returning to **Active Mode** (**AM**: 160μA, radio in RX: 15mA, radio in TX: ~30mA depending on transmit power) through enabled interrupts in less than 6μs [28, 29]. The state of execution is saved on the stack and is restored unless altered during the interrupt service routine. State and memory are maintained and, as a result, the network is able to store data while nodes are put to sleep. This was an important requirement given that we had intermittent link failures.

The network protocol stack of WizziMote is built around the DASH7 standard [30]. The data communication is simplex, i.e., the network interface can only operate as a transmitter or as a receiver at one time. The radio channel bandwidth is 1.74 MHz, spanning from 433.05 to 434.79 MHz. While DASH7 offers different classes of communication channels, in this project, we used the normal category [30]. It provides eight channels, each with a data-rate of 55 Kbps. In our experimental analysis with the WizziMote, the neighboring channel interference was high. As a result, we used only four non-overlapping channels.

4 NETWORK ARCHITECTURE AND PROTOCOLS

The wireless mesh network contained three types of nodes; **sensor nodes** that measured water temperatures, assembled the data into packets, and transmitted the packets to a backbone of **relay nodes**, which in turn moved the data packets to a **gateway node**, which uploaded the data to a server.

(1) **Sensor Nodes** *(Figure 1: orange, Figure 2(a))* interfaced with a string of five Campbell CS225-L temperature sensors, spaced either 5 or 10 m apart on the length of the sensor cable and anchored into the river bed. These nodes contained WizziMotes that connected to the wireless network and were the sources of the temperature data. At the heart of these nodes was an Arduino Mini Pro, which periodically polled sensors for the temperature data, formatted

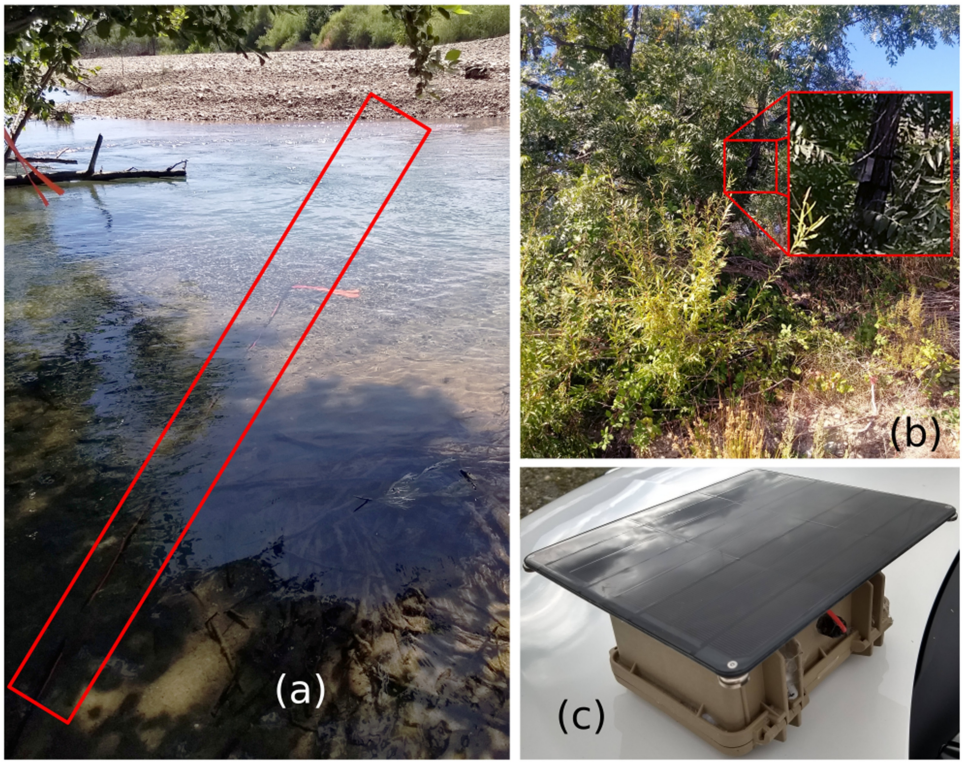

Fig. 2. Pictures illustrating the disparate nodes installed in the lower Yuba River, including a sensor node (a - red box indicates the path of the sensor), a relay node (b - inset of the relay in red), and the gateway (c).

the data into a packet, and passed the packet over a serial interface to the WizziMote that then transmitted the packet to a relay node in the wireless network. While our deployment had up to nine sensor nodes (this number fluctuated with damage and fouling by wildlife, but seven remained active through most of the sampling period), the network could handle many more sensor nodes. Each of our sensor nodes were equipped with five sensors anchored to the riverbed. The placement of sensors was based upon two criteria: (a) interest or importance of micro-habitat for juvenile salmonids and (b) feasibility of installation. Installed sensors both bisected major channels, and zig-zagged through off-channel micro-habitats.

(2) **Relay Nodes** *(Figure 1: blue, Figure 2(b))* passed data packets from relay and/or sensor nodes to other relay nodes, until the packets reached the the gateway node. These nodes contained a battery and a WizziMote. Together, these nodes formed the network of relay nodes transporting the sensor data to the gateway node. The main function of these nodes was to receive packets and transmit them. These nodes were able to buffer packets during periods when the network links failed and were unable to transmit packets towards the gateway node.

(3) **Gateway Node(s)** *(Figure 1: magenta, Figure 2(c))* was the aggregation point. On the network side, the gateway received packets via a WizziMote from one or more relay nodes. It then transmitted these data to a Raspberry P0 over a serial interface. The Raspberry Pi timestamped the data packets, stored them on a memory card, and periodically used a cellular modem to upload the data to an off-site server. In our deployment, there was a single gateway node that was powered by a solar-charged battery.

The relay and the sensor nodes were powered by non-rechargeable batteries, so energy savings was critically important for these nodes. Only the gateway node's batteries were charged by a solar panel. The most important reason for not using solar rechargeable batteries for the sensor and relay nodes was the need to hide the electronics from hikers, hunters, and animals. We needed to create a package that was small and could be hidden in trees and other hard-to-reach places where solar rechargeable batteries would not be very effective. This also meant that we could only fit a small number of batteries (sensors contained four, and relays contained only a single 19 Ah lithium thionyl-chloride batteries). This limited number of batteries further necessitated power savings. The gateway node with the solar panel was deployed in a tree and appropriately camouflaged.

4.1 The Deployed System

The network was deployed in November 2018 on the 3 km study reach of the lower Yuba River. The sensor network remained functional through May 2019. The study reach was immediately downstream from Daguerre Point Dam (Figure 1: not featured, just upstream of the top-right of the map) near Marysville, in California's Central Valley. The river is a multi-threaded river. The primary channel of the river is approximately 100 m wide in the study reach, with a bankful flow of roughly 141.6 m^3/s. The river has a course riverbed of large cobble, allowing for considerable hyporheic flows through the gravel. The river is heavily modified due to extensive aggregate gold mining in the region. This reach of river was in fact re-located a century ago, moved around the gold-fields and to the dam. As a consequence, the river is quite treacherous—surrounded by large gravel mounds. Further, the movement of the river was relatively recent, so there has been little ecological succession that would yield sedimentary organization. Indeed, much of the floodplain contains large, exposed cobble. It is a very difficult area to navigate, and low sediment organization, substantial hyporheic flows, high velocities, large cobble size, and forced channeling all contribute to a particularly dynamic study reach, prone to re-organization and movement during flood events (Figure 3). As such, sensors were placed within the river in places that were accessible by foot, without the need of a boat. Sensors were placed, orthogonal to the direction of flow, with effort made to run them through as many different geomorphic features as possible. Because of the limitations of this dynamic system, sensor placement was not random (though random number generators were used to select approximate location), field crew safety required adaptation to place sensors in areas that could be reached without workers being harmed or carried away by substantial flows.

4.2 Packet Types and Structure

There were three types of packets: **data (D; presently temperature data), health (H)**, and **test (T)**. Temperature packets consisted of a node ID (the unique identifier of the sensor node from which the data originated), sequence number (a counter that incremented at the sensor node for each transmitted packet so duplicate packets could be identified), sensor type (an ASCII character to characterize the sensor—presently "T" for temperature), message type "D" for data, first sensor ID, last sensor ID (sensor nodes were designed with multiple sensors per node—presently five—using the first and last sensor ID, we could infer whether sensor readings were ascending or descending in their ordering and allowed for the pinpointing of malfunctioning sensors), up to 20 bytes of data (in this application, up to 10 temperature readings), and two bytes of 0xFF to signify end of message.

A health packet (H) contained only metadata: node ID, sequence number, sensor type of "R" for recovery (sent when a node restarted or regained power after a power loss), and message type of "H" for health. Sent at one-hour intervals, the reception of health packets ensured that the network was functional between the originating nodes and the gateway. These packets were particularly useful during periods when network links failed. All nodes (including relay nodes)

Fig. 3. Image of a flood event in the lower Yuba River side channel, within the study reach, just downstream of the Gateway node. Marked in the image is the approximate location of a sensor (red), and a relay node (orange).

generated health packets, which allowed us to localize failures to a specific node or region of the network. A **testing packet (T)** was used for debugging purposes. Testing packets had a structure of node ID, sequence number, sensor type "T," and data type "T"; they were generated by handheld nodes, carried by researchers to test wireless communications between nodes as they were being deployed in the field, ensuring optimal node positioning and placement.

The maximum packet size was 28 bytes, set statically by the sensor node when creating packets. The first four of these 28 bytes were reserved for the packet metadata. The actual packet could be smaller, depending upon the amount of data collected by the sensors connected to the sensor node.

4.3 Medium Access Control (MAC) Protocol

Due to the limited access to the lower Yuba River, the lack of access to power, and the threat of vandalism posed to visible structures like solar panels, it was imperative to minimize power use. The sensor nodes aggregated temperature data from the sensors and transmitted to the next upstream relay node every 15 min. This implied that all sensor nodes could be powered down to low-power states with a duty cycle of 15 min. A related issue was that nodes that transferred data to the same upstream node use the same frequency. This led to collisions, which needed to be prevented to minimize power use.

Energy-efficient MAC protocols have been extensively studied in the context of wireless sensor networks with a sleep-awake duty cycle. Broadly, they can be categorized into four classes: synchronous, asynchronous, hybrid, and random [31]. Due to the lack of real-time synchronized clocks in the nodes, synchronous approaches such as T-MAC [32] and hybrid approaches that combine synchronous and asynchronous approaches could not be used. It was further impractical to implement any wireless time correction like GPS, as GPS reception can be impacted by vegetation, and this interference can yield errors in both positional and timing signals [33]. Additionally, GPS-based timers used considerably more power than was available in our installation. GPS receivers were also more expensive than our budget allowed. With low duty cycle sensors (such as those

implemented in this project), it has been found that idle listening can be reduced by having the transmitting node be the active entity when synchronizing with the receiving node [31]. A common way for this to be achieved is for the sender to—when it has data—transmit a preaamble that lasts for a length of time that is long enough to overlap with the short duration of time for which the receiving node wakes to listen for the preamble. This overlap is guaranteed even if the clocks in the transmitting and receiving nodes are not synchronized. If the receiving node detects the preamble, then it stays awake to receive the data. This is the approach adopted in B-MAC with additional optimization in X-MAC [31]. Recently, a randomized algorithm based on the Birthday protocol [34] has also been implemented [35].

The MAC protocol adopted by WizziMote was similar to B-MAC [36]. In the WizziMote, time was measured in "ticks," with each tick equal to $1/1,024$ seconds. For our deployment, we set the sleep duration to 900 ticks, while any advertisements were transmitted for 1,000 ticks. This ensured that receiving nodes woke up before the end of the advertisement period of the transmitting node. During the 1,000 ticks, the transmitting node continuously sent packets containing the remaining time of advertisement (in ticks). The receiving node used this information to synchronize with the start of the data transmission. At the transmitting node, the packet was transmitted immediately after the completion of the advertisement. Once the data were transmitted, the transmitting node switched the radio interface to the receive mode to receive the acknowledgment. At the receiving node, if the data passed the checksum, then the node immediately acknowledged the transmission on the same (backward) channel. Notably, the acknowledgments did not use advertisements. The data transmission and the acknowledgment were performed in one single operation [37]. Furthermore, nodes transmitted data regardless of whether or not any node received the advertisement due to a link failure. If the transmitting node did not receive an acknowledgment, then it reattempted the transmission in between sleeping until the receiving node acknowledged the data packet. The reattempts were guided by an exponential back-off algorithm (Section 4.5).

The energy savings of this approach was derived from the fact that the sender advertised only when it had data to send and the receiver was only awake for only a small fraction of the time. The use of advertisement to synchronize nodes significantly increased the chance of collision at locations where multiple branches of the network converged (i.e., where multiple nodes transmit to a single node). At these locations, multiple nodes could be advertising, transmitting, or acknowledging at the same time. This yielded considerable interference in laboratory experiments. These transmission collisions substantially increased transmission times and diminished network fidelity, filling the distributed network storage (discussed in Section 4.5). To avoid these collisions, while still preserving the advertisement-based synchronization, we implemented cross-listening: The transmitting node listened on the channel before it started to advertise on the channel. This was similar to **Carrier Sense Multiple Access (CSMA)**. The main difference was that traditional CSMA with **collision detect (CD)** and CSMA with **collision avoidance (CA)** required the radio interface to operate in full-duplex mode, i.e., able to transmit and receive at the same time; whereas the WizziMote radio interface operated in simplex mode. Thus, in cross-listen, the transmitting node only listened for the advertisement and not a general carrier signal before any data transmission. By minimizing channel conflicts, reducing contention, and decreasing transmission attempts, cross-listening also reduced duplicate packet transmissions, decreased the time for a packet to reach the gateway node, and reduced network-wide power use.

4.4 Routing Protocol

Routing protocols in the context of wireless sensor networks have been extensively studied [38] and there is a standard routing protocol defined for low-power, lossy networks like wireless sensor networks [39]. There were multiple constraints that limited the design and implementation of the

routing protocol. First, WizziMotes did not have any in-built node discovery and routing algorithm. As a result, any routing algorithm had to be implemented from scratch. Besides the complexity of many of the standard algorithms, there was limited memory in each WizziMote (see Section 3.3). Additionally, there were both topographical constraints (dense foliage and sharp embankments at many locations) and constraints on the number of potential locations where relay nodes could be deployed. Consequently, the network of the relay nodes was a sparse network and there were only a few routes between the sensor nodes and the gateway node. Near many of the sensor nodes there was only a single reachable relay node.

There were a number of conflicting constraints imposed upon our network; in balancing them, we opted for simple static routing and implemented network storage to account for temporary link failures. The locations of the temperature sensors were fixed by the scientific requirements. Specifically, we needed to measure temperature along a 3 km river reach at regularly spaced locations. At each location, we needed multiple, evenly spaced (5–10 m) sensors. The sensor nodes that aggregated the cross-channel temperature at different locations along the river were constrained by the river's inherent geometry. The harsh terrain prevented a very dense deployments of relay nodes and was further complicated by the incidence of fouling and vandalism. The equipment and deployment budget limited the number of relay nodes. We therefore deployed a relatively sparse network of relay nodes. Furthermore, the 433 MHz radio that was implemented in the WizziMotes could reliably discriminate only four out of the seven available channels. For all of these reasons, we implemented a simple static routing strategy and used network storage to mitigate link failures.

To simplify the routing of packets and reduce energy use by eliminating the need for node discovery and distributed topology creation, all routes in the network were set statically. To minimize interference, nodes that were in the hearing range of each other but on different routes were assigned different channels. The channels were reused in links that were far apart. Developing the network with static routing simplified the implementation and provided greater control over the exact traversal path of packets. This came at the cost that there were no alternate paths for packets when field conditions changed. But the limited deployment meant that in most places on the river, no alternative path existed.

A potential approach to deal with link failures would be to set up static back-up paths. If the primary path broke, then the back-up path could take effect to maintain the connectivity of the network. However, this option would only be viable with a dense deployment of relay nodes. In fact, for a dense deployment of relay nodes and with more power, processing, and memory in the relay nodes, a dynamic routing protocol would be beneficial. The overall topology of the sensor and the relay nodes had an unbalanced tree structure with the gateway at the root node (Figure 1). It was only at merge points that nodes in tree branches were close enough to hear each other. We used cross-listen at these points to reduce collisions.

4.5 Network Storage

By doing a pilot study with a small, experimental field deployment, we found that weather, distance, local topography, vegetation, antenna position and direction, and other factors reduced link fidelity and intermittently hindered transmissions. Further, transmission were lost to collisions at points where network branches merged. For times of diminished network stability, we designed modules that minimized repeated re-transmissions while maximizing the throughput of unique packets. This was achieved by exponentially backing off of the transmitter when transmissions failed. For example, if the base transmission rate was every 4 s, then after subsequent unsuccessful transmissions, the base was binary exponentially increased to 8 s, 16 s, 32 s, and so on. When the link quality was impacted by weather, it often remained so for some time. Thus, the exponential back-off (as opposed to a linear back-off) reduced the number of unsuccessful transmission

attempts, and saved power. While the node slowed down its transmissions, it still received packets until its queue was filled. This, in turn stored as many packets as possible within the relay nodes in the network. Once the link became reliable—or in other words, once the node successfully transmitted a packet again—it quickly recovered from exponential back-off by immediately resetting its base to four seconds.

We have separated the network storage algorithm into two parts (**Algorithm 1**): The first was done by the transmit thread, wherein it retrieved the number of unsuccessful packet transmissions, and then the receive thread updated the back-off timer. This back-off timer reduced the frequency with which transmissions were attempted. In our current implementation, the delay began at four seconds between transmission attempts and doubled for each transmission failure, up to a maximum of 30 minutes. This maximum ensured that each node attempted a transmission at least every half-hour. Any successful transmissions restored the delay to four seconds.

ALGORITHM 1: Network Storage Algorithm

Input: Back-off timer T_B; Counter C_f for failed transmission in the last round; Transmission base value B_B; Max transmission back-off B_{MAX}; Minimum transmission back-off B_{MIN}.

```
1: procedure TX(C_f, B_B, T_B)                          ▷ For any TX round i
2:     if Successful transmission then
3:         C_f ← 0
4:         B_B ← T_MIN                                  ▷ reset transmission base
5:         T_B ← 1
6:     else if Unsuccessful transmission then
7:         C_f ← C_f + 1
8:     end if
9: end procedure
```

```
10: procedure RX(C_f, B_B, T_B)                         ▷ For RX round i+1
11:     if C_f > 0 and B_B < B_MAX then
12:         C_f ← 0
13:         B_B ← B_B × 2
14:     end if
15:     if T_B mod B_B == 0 then
16:         T_B ← 1
17:         call TX procedure
18:     end if
19: end procedure
```

These transmission delays coupled with the 10-slot queues in each relay node coalesced to yield up to a few hours of distributed network storage (depending upon the design and topology of the network). This data retention was crucial for scientific sensing, which could generate a steady stream of sensor-derived data packets. Their timestamps were reverse-calculated once they successfully reached the gateway (Section 6).

5 NODE SOFTWARE ARCHITECTURE

In discussing the software architecture of the nodes, we focus on the relay nodes, as they were the only nodes that contained both transmit and receive functions.

5.1 Event-driven Proto-threads

Nodes in the network utilized the WizziMote's **Kernal Abstraction Layer (KAL)**, which maintained the illusion of multi-threading at a high level by encapsulating the Contiki proto-thread library [27], an open-source operating system that offers dynamic loading of application code onto embedded systems. While similar to a threaded architecture, pseudo-threads cannot run in parallel. Proto-threads are lightweight and stack-less, providing a blocking context on top of an event-driven kernel. Sequential flow of control is implemented with macros that save the processing states of functions without using complex state machines or full multi-threading. Proto-threads were particularly well suited for memory constrained devices (WizziMotes had only 4 KB of memory). Another useful feature is that preemptive multi-threading is implemented in the proto-threads as an application library. Consequently, multi-threading can be optionally linked with programs that explicitly require it [40].

Two types of events—timer and radio—controlled a pseudo-thread's execution flow. A pseudo-thread was always in one of the following states: inactive, processing an event, or waiting for an event. The WizziMote radio library was written with only one radio buffer. While both the transmit and receive threads could be active at the same time, they could not process radio events simultaneously because of the shared buffer. To prevent conflicts over radio resources, these threads were scheduled sequentially (transmit after receive).

5.2 Packet Queue

A finite packet buffer, implemented as a queue, connected the threads and acted as shared storage. Packets received from the receive thread were inserted into the queue, which were then removed from front of the queue by the transmit thread. To reduce network load from re-transmissions that resulted from unheard acknowledgments, the receive thread compared each received packet with the packet received before it. Identical, sequential packets were not inserted into the queue.

5.3 Thread Architecture

In each relay node, there were three threads: the management thread, the **transmit (TX)** thread, and the **receive (RX)** thread (Figure 4). At points of intersection in the topology, we also introduced the cross-listen thread for collision control.

5.3.1 Management Thread. (**Finite State Machine (FSM)**: Figure 4) as the name suggests, managed the transmit and receive threads by initializing them with the necessary parameters. There was also a main() function that initialized the **Hardware Abstraction Layer (HAL)**, KAL, and the management thread. The main() function and management thread coexisted because all KAL processes had to exit before the node could transition into sleep mode. Since the management pseudo-thread controlled the other threads, killing the management thread terminated the other threads, which yielded a clean state.

In addition, the main() function controlled two timers: the main timer, which put the node to sleep (see Section 5.4), and the watchdog timer, which detected inactivity. The MSP430's built-in watchdog timer performed a hardware restart if the timer was not reset within 16 seconds, which allowed for recovery from any hangs or bugs. This was particularly useful in our implementation, as elements of the WizziMote's software were provided pre-compiled and we did not have access to their source [28].

5.3.2 Receive (RX) Thread. (FSM: Figure 5) listened on the backwards channel for any incoming messages and inserted them into the queue. If nothing was on the channel, then the node switched to the transmit thread to transmit any packets that were in the queue. If the RX thread successfully

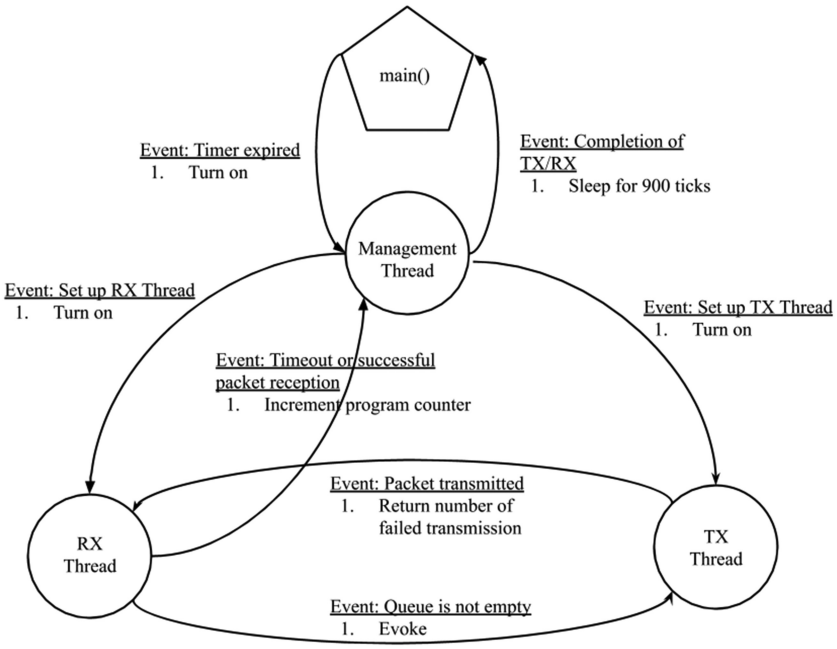

Fig. 4. Finite State Machine (FSM) showing the overall organization of the pseudo-threads and the events that triggered transitions to the different threads.

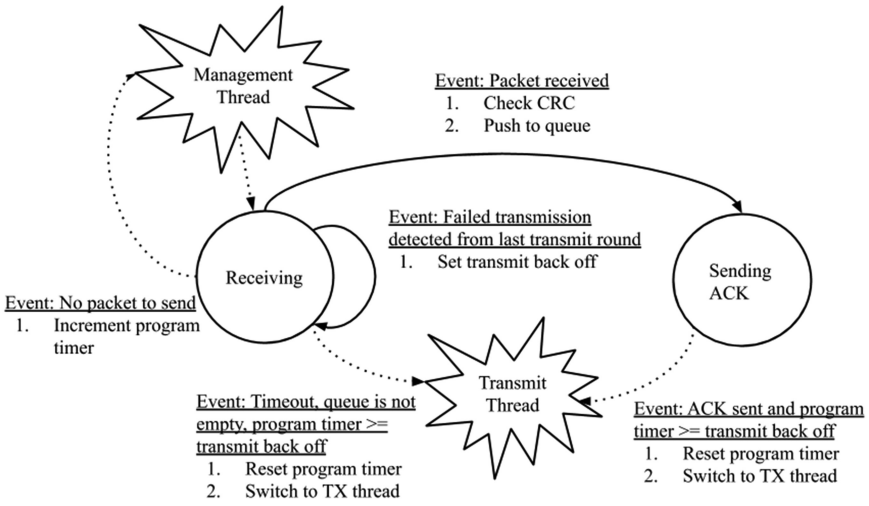

Fig. 5. The Finite State Machine of the RX thread referenced in Figure 4.

received a packet, then the thread acknowledged the packet to the transmitting node on the backwards channel, which signaled a successful transfer.

In the sequence of events upon waking up from sleep, the receive thread was instantiated before the transmit thread. This was to fill the queue as much as possible to maximize the storage utilized in the network. Even when the forward channel was broken, the node continually tried to receive until the queue was full.

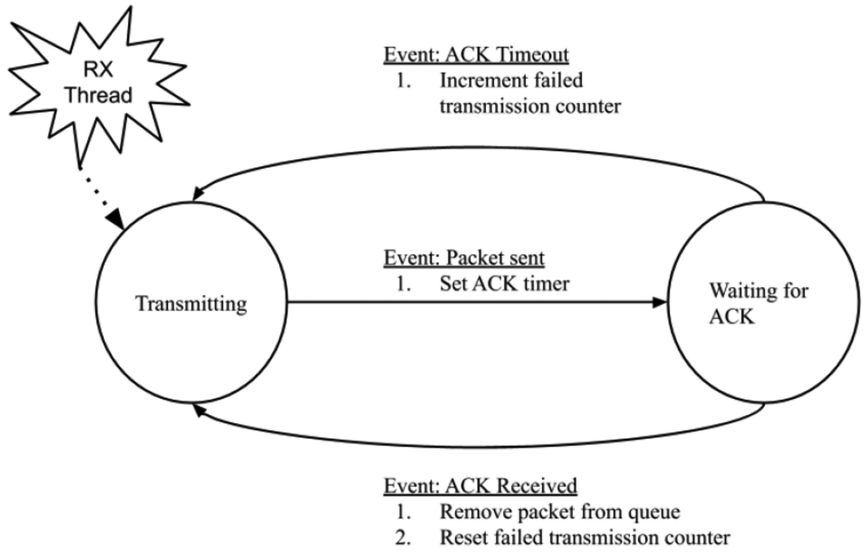

Fig. 6. The Finite State Machine of the TX thread referenced in Figure 4.

5.3.3 Transmit (TX) Thread. (FSM: Figure 6) transmitted packets (if there were any in the packet queue) on the forward channel. First, an advertisement was transmitted continuously on the forward channel to signal that a transmission was to be sent. The advertisement period had to be longer than the sleep period of the receiving node to ensure that the receiving node's listening period would overlap with parts of the advertisement. At the end of the advertisement, the node transmitted the data packet, regardless of whether a receiving node was active to listen for it or not. Finally, the transmitting node listened for an acknowledgment, which if received, prompted the removal of the transmitted packet from the sending node's queue. To minimize power use, the transmit process was not invoked every cycle of the node. If the forwarding channel was down, then the node only transmitted after the delay timer had expired (Section 4).

5.3.4 Cross-listen Thread. There were nodes at confluences within the network, where two or more nodes transmitted on the same channel to the same receiving node. In these nodes, an additional thread was added to minimize transmission collisions. Acting much like the receive thread, the cross-listen thread listened on the forward channel to see if another node was already transmitting. If it heard no transmissions, then it notified the TX thread to transmit the packet. If the channel was in use, then it backed off for a period of time. This delay was randomized to prevent additional collisions that could propagate from non-random or hard-coded delays.

5.4 Energy Efficiency

For nodes to survive in the field with limited battery power, they were put to sleep when not transmitting or receiving. During this sleep, each relay node was pulled to Low Power Mode 3 (LPM3), which disabled the **CPU (MCLK)** and the **high frequency clock (SMCLK)**, leaving only the 32 kHz low-frequency crystal clock ACLK powered [29]. In the main() of each relay node, a timer of 900 ticks was initialized and started, which created an interrupt and returned the node to Active Mode.

```
main_timer = kal_timer_start(KAL_ETIMER, NULL, TIMER_ID, 900);
```

A few lines later, after closing the radio layer and the serial, the node was put into LPM3 and global interrupts were enabled. The node woke up when the main timer timed out.

```
_BIS_SR(LPM3_bits + GIE);
```

This process was repeated after every cycle of receive and transmit.

6 DEPLOYMENT AND EVALUATION

The wireless sensor network was installed and active from November 2018 through May 2019. Many of the actual sensors were secured into the river over the summer of 2018, as low flows exposed more of the riverbed. The networking hardware was added to the sensors in October and November 2018 and brought online as it was completed. The full network was completed and functional in January 2019. During the initial deployment, some relays needed to be re-positioned or relocated to address changes in flow and vegetation (Figure 1). In May 2019, a restoration project began, which substantially modified the river channel in the study reach. This work rendered network maintenance impossible. In January and February 2019, much of the river became inaccessible due to flooding, and repairs were more challenging. The low-power, small footprint of all hardware yielded a relatively low-effort field deployment. Field trips were made as day trips from UC Davis, often with just a single worker in the field. Additionally, because all hardware was programmed in advance and sealed prior to the field deployment, field workers needed no specialized skills.

6.1 Power Use

To determine the power used by our relay nodes, we conducted an experiment that logged the current drawn by relay nodes. We did this by setting up three nodes: a dummy-sensor node that generated numerous packets at regular intervals, a relay that transmitted the data, and a gateway node that received the data from the relay. We set up the dummy-sensor node to transmit at regular intervals of 4, 8, 16, or 32 seconds. The relay node was connected to a 3.600 VDC source, through a high-speed, logging bench-top multimeter to measure draw. We measured voltages both above and below this logging meter to ensure that burden voltage was not too large and that supply voltage to the relay node stayed above 3.500 VDC. These voltage monitoring multimeters were both put into their high-speed max/min modes, which measured at least once every 1 ms. These multimeters were recently calibrated; one was an Amprobe AM-270 and the other a Fluke 87. The power source was a recently calibrated Power Designs 4010. The current was logged by a recently calibrated Rigol DM3051. This meter was capable of 500 samples per second, with storage for 200,000 samples—yielding 6 minutes and 40 seconds of continuous sampling.

Each of the relay node's states were reflected consistently in the data (Figures 7 and 8). Low-power sleep states used 2–3 μA, active mode used about 5 mA, and receive and transmit modes used 20–30 mA (Figures 7 and 8). These values were consistent with the anticipated power use from the CC430F5137 datasheet (Section 3.3) [28, 29].

6.2 Network Storage

To test the effectiveness of the network storage, we created a simple experimental network in the laboratory that simulated link instability and observed the network recovery and loss. We considered a linear network of a sensor node connected to a gateway node through two relay nodes. The sensor node transmitted data every 30 seconds, and the relay nodes had a packet queue of 10 packets. The link between the second relay and the gateway was shut down with increasing

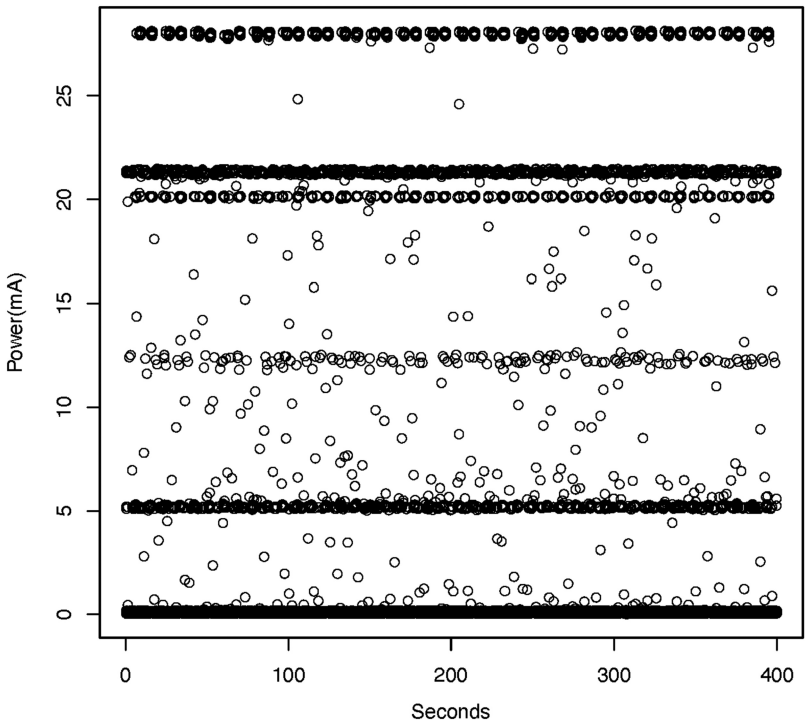

Fig. 7. Raw power-use data taken from the Rigol DM3051 in mA when relay node received a data packet every eight seconds. It is notable that the different power modes were very consistent. The readings near 0 mA were low-power sleep states (2–3 μA), the 5 mA readings corresponded to active mode, and those in the 20–30 mA range corresponded to receiving and transmitting. All others are transitional between states.

duration of 10, 100, 600, and 1,000 seconds. Each link failure was followed by 60 minutes of potential recovery time. Recall that the definition of network recovery in our network storage algorithm is equivalent to a successful transmission through the broken link (**Algorithm 1**). Hence, we defined the time for network recovery time as the time difference between the link failure and the receiving of the first packet after the link became available again. Note that there existed an idle time period during which the network has physically recovered, but was undiscovered until a packet traversed the previously broken link. For the purpose of this application, we considered the link to be dysfunctional during this idle time period, because the network health was dependent on the overall traversal of packets in the network rather than the individual link health.

Figure 9 shows the box plot of the recovery time for different durations of link failure. As the link breakage time increased, the time to recovery also increased. As the link was broken for longer periods of time, the delays in between transmissions increased exponentially and so did the recovery time of each node. The queue size also influenced the back-off rate of the network. When the queues of connected relays filled, the overall recovery time of the network increased as multiple relay nodes' back-off durations compounded.

The rate of packet creation was significantly greater than in a real-world deployment, which ensured that the packet queues would fill while the link was broken (particularly during the longer duration link failures). The purpose was to test the back-propagation of packet storage. This effect was observed by looking at the losses. For any given queue size, breaking any link in the

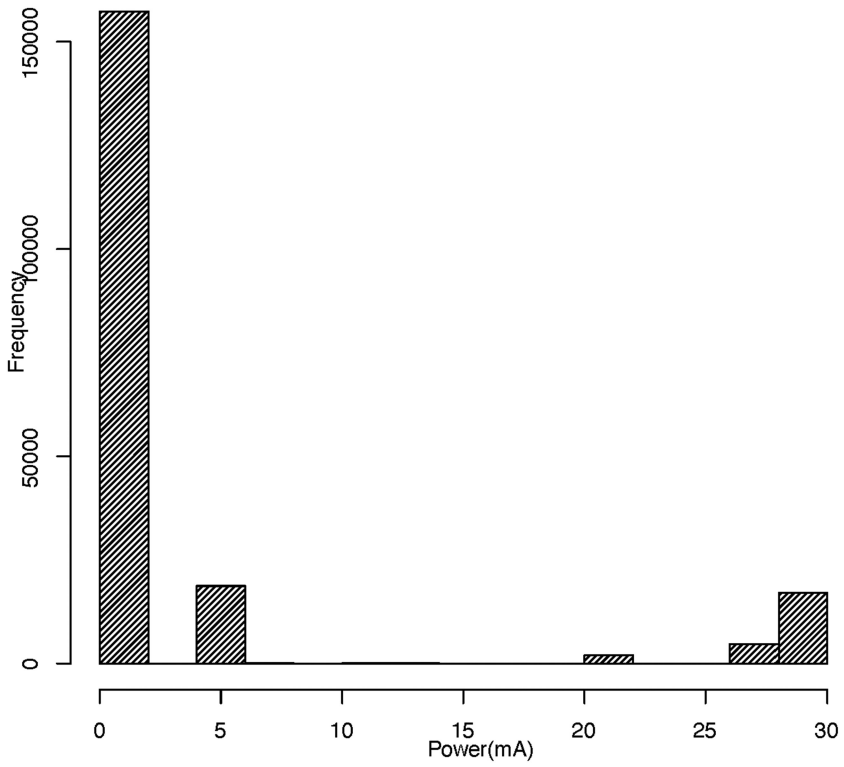

Fig. 8. Histogram of the power used in mA when relay node received a data packet every eight seconds. Each of the peaks in this histogram corresponded to a state for the Wizzimote. The peak around 0 mA corresponds to the low-power sleep state, 5 mA corresponds to active mode, and the peaks between 20 and 30 mA correspond to transmit and receive modes.

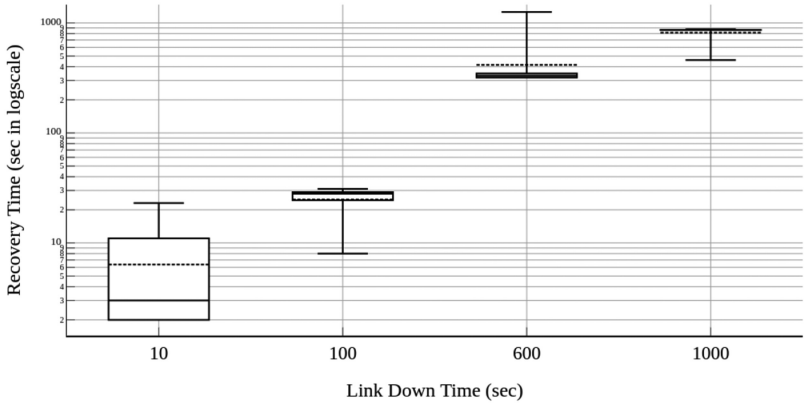

Fig. 9. Network recovery time in seconds (log scale) as a function of link failure duration. The mean recovery times are shown as the dashed line and were, respectively, 6.36 s, 24.83 s, 415.36 s, and 815.75 s. The median recovery times were, respectively, 3 s, 28 s, 331 s, 865.5 s.

network had no effect on the packet loss if the down time was relatively low. Once the queues of all available relays were filled, the packet loss increased in proportion to the link down time. In a real-world implementation, the packet frequency would be much smaller, and the number of relays much larger, granting the network more storage over a longer duration of time, minimizing the probability of packet loss.

6.3 Timestamp Alignment

The distributed network storage implemented in this network resulted in the accumulation of unknown delays during the transit of sensed temperature data. The network contained no real-time clocks, and packets were timestamped upon receipt at the gateway, after any applicable network delays. It was necessary to re-align these packets with real-time. Re-alignment was feasible, as each packet contained an 8-bit packet ID, which was sequential (rolling over at 255). We performed the re-alignment in a two-step process, first identifying improperly timestamped packets (Figure 10(a)), then re-aligning them (Figure 10(b)). The central idea underlying both the identification and re-alignment of packets was the use of neighboring packets to estimate the timestamp for each packet.

While packets from other sensors may be useful for aligning packets in situations where transmission delays propagate from sensor or network drop-outs, most of our network failures were the result of power-failures at the gateway. Since timestamps were generated at the gateway, and most link failures propagated from the gateway (as it was the only solar-powered unit in the network and was most prone to power loss during inclement weather), we did not use packet timestamps from other sensors to align packets, opting to instead use packets from each sensor to exclusively align timestamps from that same sensor.

Each sensor generated a data packet about every 15 minutes, and each packet contained a sequential packet ID. It was therefore possible to estimate the timestamp of any individual **packet (P)** using the timestamps of neighboring packets (X of length N). By taking the difference in packet ID and multiplying by time between sensor readings (i.e., 15 min), we were able to estimate which packets were properly aligned. The method was robust to many different modes of packet delay, because it treated most aspects of packet timings as variable. To this end, we knew that the time between sensor readings (T) was determined by a crystal, and never synced to any clock, and therefore fluctuated with temperature. To improve timestamp estimates, we began by setting T=14 min, then increased it in five-second increments until the model heuristics (Figure 10) indicated that we had reached an optimal timing for each packet. When evaluating whether timestamps were aligned (Figure 10(a)), we evaluated the range of estimates, the distance of the estimate from the observed packet timestamp, and the number of estimates used to generate the estimated timestamp. In each step, we also removed the minimum and maximum estimate in an effort to improve the mean estimate and decrease the spread of estimates. If the estimates derived from a set of neighboring packets failed to meet the appropriate model heuristics, or if after filtering out minimum and maximum, the number of estimates fell below L (L=6, presently), then we increased the number of packets considered (N) to include additional packets before and after the packet in question. In this way, we prioritized using the packets closest to the packet in question, expanding the window only when estimates were not sufficiently robust using a smaller numbers of packets. If N reached 72 (36 packets before and after P), then we marked packet P as bad, meaning that we could not create a sufficiently robust estimate of the correct timestamp for P. Once each packet P had been evaluated, any marked bad were removed from the data and the algorithm started over, evaluating all remaining packets not identified as bad in previous rounds and using only the remaining packets not already marked bad to evaluate each packet within the reduced dataset. This repeated until the list of bad timestamps grew by less than 5% in a given round. At this point, we considered

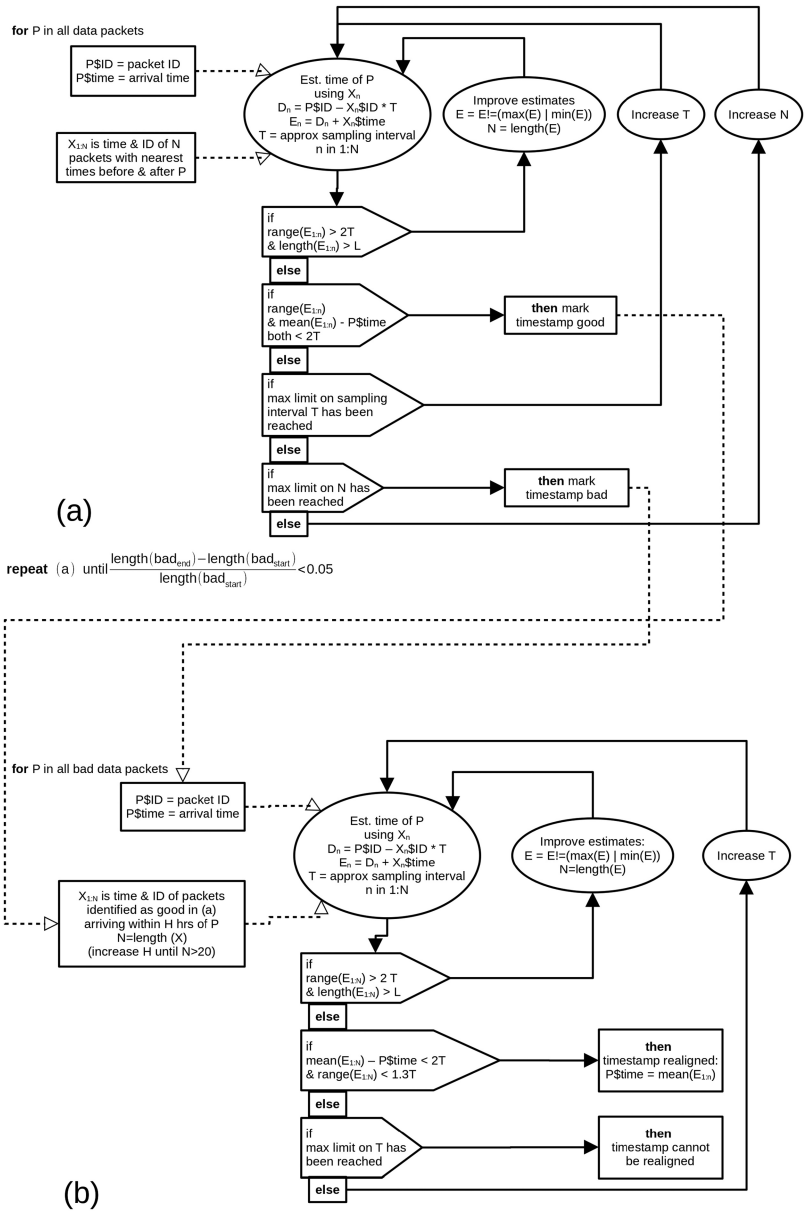

Fig. 10. Flow of control (filled arrows) and data (empty arrows) for (a) bad timestamp identification and (b) bad timestamp re-alignment. In both diagrams, T is the approximate time between data packets sent by the sensor (about 15 min). P is the packet that is being evaluated or re-aligned, X are the N-packets that are being used to evaluate the timestamp of P (notably, the selection of these packets differs between (a) and (b)). L is a means of ensuring that despite the removal of outlying time estimates, sufficient estimates are retained to ensure robust estimates.

the list of bad packets to have been identified for that sensor. These bad timestamps were then fed into a timestamp re-alignment algorithm (Figure 10(b)).

The process for timestamp re-alignment was similar to that of bad timestamp identification. We used neighboring packets that were not identified as bad (Figure 10(b)). Instead of using the

Table 1. Number of Timestamps, Broken Out by Sensor, Including Those with Bad Timestamps as Well as Those that Were Able to Be Realigned

Sensor	Date of First Packet	Date of Last Packet	Total Packets	Bad Timestamps	Corrected Timestamps
A	2018-12-05	2019-02-13	3,578	710	288
C	2018-11-06	2019-05-21	7,247	1,850	487
D	2018-11-12	2019-05-21	8,596	2,241	1,032
E	2018-11-01	2019-05-25	12,205	1,185	764
F	2018-11-07	2019-03-02	6,446	866	312
J	2018-12-14	2019-02-23	4,073	1,592	630
K	2018-12-14	2019-03-06	3,571	1,150	377
Total	2018-11-01	2019-05-21	45,716	9,594	3,890

Total number of packets is the sum of good and bad timestamped packets. Re-aligned timestamps are a subset of bad timestamps—all others could not be re-aligned.

N-nearest packets as above, we instead used good packets arriving within H hours of P. H was initialized at 7 hours before and after the documented timestamp of P, increasing iteratively until the **number (N)** of **packets (X)** used to estimate the timestamp of packet P was at least 20, with H having a maximum of 48 hours. Once H was set and the **N-packets (X)** within H hours of P, the timestamp of P was estimated and the minimum and maximum estimates were eliminated until the range of estimates was less than 30 min. If this could not be reached, then as above, the inter-packet **timing (T)** was incremented by 5 seconds and the process was repeated. This was repeated until T reached 16 minutes.

The parameterization of variables and coefficients in our implementation were largely based upon the scientific questions at play. We were interested in hourly river temperature data. As such, we sought at least two measurements per hour. This is why we chose to set our rejection thresholds at 30 min for both for evaluating whether or not timestamps were correct, relative to timestamp estimates (i.e., mean (timestamp estimates) - observed timestamp) as well as ranges of timestamp estimates. Other applications of these methods should use parameters informed by relevant research questions.

We identified 9,594 incorrectly timestamped packets. Of these, 3,890, or about 40%, were able to be re-aligned using this procedure (Table 1). As a result of this identification and realignment, 40,012 packets were able to be used in further analyses of fluvial water temperatures. Notably, each packet contained five sensor readings, yielding a total of 200,060 properly aligned or re-aligned temperature readings from the sampling period (Table 1). For most sensors, timestamp re-alignment was most common for network delays of under 10 hours (Figure 11). Indeed, for every sensor except sensor E, 75% of re-aligned timestamps were shifted less than 10 hours (Figure 11). However, for sensor E, which was closest to the gateway, large delays were more readily fixed during timestamp re-alignment. As a result, the median re-alignment was much larger (over 10 hrs: Figure 11). Notably, timestamp re-alignment was feasible for up to 20 hours for at least some packets from each sensor (Figure 11). Timestamp alignment utilized the sequential packet ID, which rolled over at 255. Any timestamp alignment was infeasible after 63.75 hr, when the counter rolled over. Re-alignment also successfully corrected hourly mean temperatures, bringing them in line with what would be expected, with lowest temperatures in the early morning and warmest temperatures in the late afternoon (Figure 12). Prior to timestamp processing, the mean temperature was elevated in the early morning. Timestamp alignment rectified these high mean temperatures, as these incorrectly timestamped packets were shifted to the correct time.

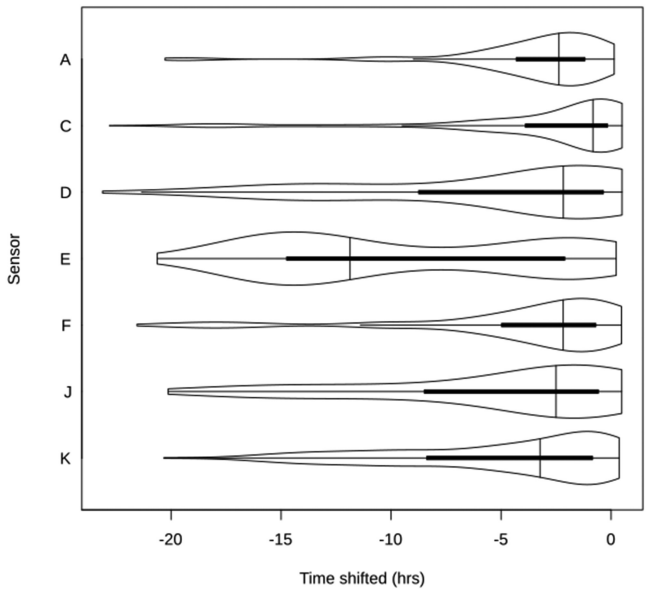

Fig. 11. Violin plots of the amount of time shifted during timestamp re-alignment, broken out by sensor. Violin plots also contain box plot, with whiskers at 0th and 100th quantile, and box extents to 25th and 75th quantiles. Perpendicular lines are at medians.

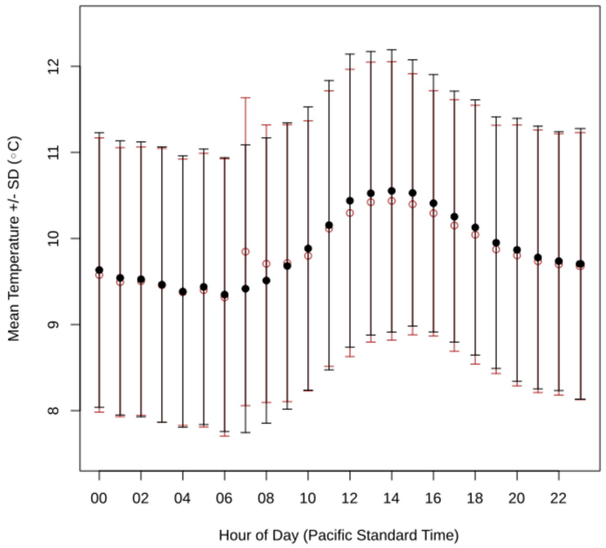

Fig. 12. Temperature data, broken out by time of day (local, Pacific Standard Time), with mean and standard deviations in temperature across all sensors for the whole of the sampling period. Red, unfilled circles demonstrate the data as collected. The red plot includes the data that were rejected during timestamp alignment and also uses the raw, unprocessed timestamps. The black, filled circles are the data after processing: All incorrectly timestamped packets have either been re-aligned (if feasible) or removed.

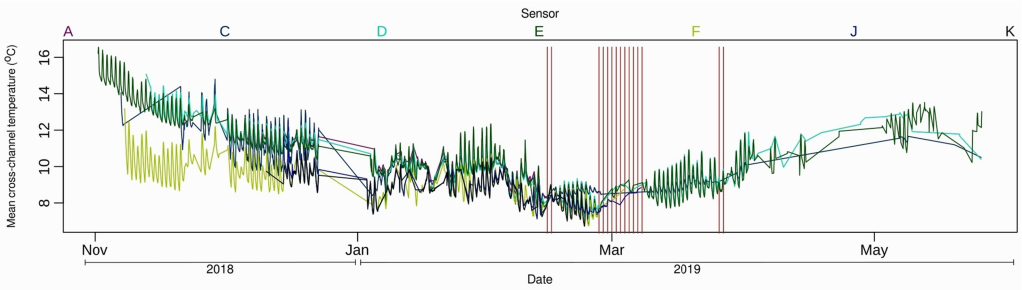

Fig. 13. Temperature data collected from the wireless sensor network. These temperatures represent cross-channel means: They are the average of the five individual sensors attached to each sensor-node. The vertical red lines are at midnight on days in which there was a flood of at least twice ($283m^3/s$) the bankful flow ($142m^3/s$).

6.4 Temperature Data

We collected temperature data from November 2018 through May of 2019 (Figure 13). In this time, we collected 200,060 temperature measurements. Different sensors intermittently failed at different times, as a result of varying network failures (Figure 13). Many were damaged during large flow events (Figure 13: red lines). Notably, sensor E persisted longest, as it was nearest to the gateway and subjected to the fewest relay failures (Table 1). Upon re-aligning the data, we found a predictable temperature curve, with hourly mean temperatures peaking in the late afternoon and reaching a minimum in the early morning (Figure 12).

In total, 3,890 packets were re-aligned. We initially questioned the balance between the uncertainty implicitly introduced by timestamp re-alignment and the value of additional data. We came to the conclusion that re-alignment was crucial, because while it represented a relatively small number of packets (only about 11% of the total packets found to have been properly aligned: Figure 10(a)), timestamp mis-alignment was biased towards packets collected at night or during inclement weather, when network failures were most common. Eliminating 3,890 sensor readings explicitly collected at night and during rain would have hampered our ability to model river temperatures during these important times. This re-alignment was therefore crucial to the success of the network and the underlying ecological science (Figure 12).

7 LESSONS LEARNED

This project began in 2014 and spanned five years from its conception through the design, deployment, data collection, data post-processing, and analysis. It provided a unique opportunity to develop and deploy a specialized wireless sensor networks for scientific applications in harsh conditions. The following are some of the important lessons learned from this project:

(1) Sensors developed to measure temperature across the channel were implemented in a thick cable, which was anchored to the riverbed (Figure 2; Section 3). Despite these steps, some of the sensors were damaged during flood events or by wildlife. More robust sensor cables could have reduced the hazard rate and extended the service life of the sensors.
(2) Even though the gateway was powered by solar rechargeable batteries, it failed during extended dark periods such as nights followed by cloudy or rainy days. To increase resilience to dark periods, the gateway could have been equipped with additional batteries.
(3) The MAC protocol that was implemented in the WizziMote was a transmitter-driven advertisement-based synchronization approach (Section 4). In such a design, the energy

cost of synchronization is predominantly borne by the transmitter. An alternative approach would be a probabilistic approach, such as that based on the birthday paradox [34, 35]. In this approach, the transmitter and the receiver randomly advertise and poll the channel and it can be shown that they have a very high probability of synchronization even with a relatively low number of polls. The energy cost of synchronization is more balanced between transmitter and receiver, though not completely balanced, as the receive and transmit costs are not the same. This approach was evaluated and found to be energy-efficient in a field deployment [35].

(4) An important lesson pertained to our reliance on the closed-source radio library. The WizziMote radio library was not open source. This hindered our ability to address bugs and develop the link layer protocol. For example, since we did not have access to the source code of the radio library, it was impossible to implement the birthday protocol [35]. Additionally, the radio only reliably discriminated 4 out of the 7 channels standardized in the DASH7 protocol stack. This also limited our channel assignment.

(5) The DASH7 protocol allows the transmit power to be dynamically adapted. Dynamically increasing the transmit power during inclement weather when the link quality deteriorated could be adapted into our back-off algorithm to save power and improve successful transmissions. While it would significantly increase the complexity of the transmit protocol, in the long run it could save energy by reducing the number of reattempts during periods of poor link quality.

(6) As mentioned in Section 6, we did not timestamp the data at the sensor nodes, because the clocks in the WizziMotes were not accurate, and there was significant clock skew. Instead, we included a sequence number in the packets. The data were timestamped at the gateway node. However, when packets were stored in the network due to link failures, the timestamps had to be aligned using the sequence numbers (as described in Section 6.3), which was prone to errors resulting in some packets to be discarded (Table 1). One heuristic would be to implement a millisecond counter of delays in the packet header as the packet proceeded from the sensor node to the gateway node. This delay could significantly improve efficacy of the timestamp alignment algorithm [35].

8 RELATED WORK

Wireless sensor networks (WSNs) [11, 12] have been and continue to be used in many different applications. These include monitoring application for early detection of forest fires [13], actuating applications in precision agriculture [14], energy usage control application in smart homes [15, 16], and tracking applications such as animal telemetry [18]. Wireless sensor networks have been used in habitat and environment monitoring [17, 41–43]. One of the earliest comprehensive works in this area was the habitat monitoring on Great Duck Island [41]. The project monitored many environmental parameters, including temperature, barometric pressure, humidity, and others. Furthermore, due to the remote location, sensor power use was an important design constraint. The paper provided significant details on the design of the sensor node, the wireless sensor network, the data aggregation through patch gateways, and power budget and evaluation. Another seminal work on the study of scalability of wireless sensor networks was the measurement study performed in GreenOrbs [42]. Based on detailed measurement study of congestion in a large deployed wireless sensor network, the study proposed novel event-based routing strategies that were appropriate for large-scale sensor networks. Another relevant study was the CitySee project [43], which used wireless sensor network for real-time monitoring of CO_2 in a large urban area. The paper addressed the challenges of deployment. Despite the similarities of these prior works, each of these environmental deployments, including the river deployment in this article, have unique features that address

specific challenges, and as a result the specific techniques and approaches were not necessarily transferable.

Techniques to minimize energy use in wireless sensor networks have been extensively studied [44]. In Reference [45], a **Sparse Topology and Energy Management (STEM)** algorithm was proposed to efficiently wake up nodes from a deep sleep state without the need for an ultra low-power radio. The key innovation was to allow tradeoffs between energy efficiency and the latency incurred to wake up the node. A different tradeoff, between data fidelity and energy efficiency, has been investigated [46]. There has also been a number of studies that have proposed energy-efficient routing schemes [47, 48].

In our work, we have used an advertisement-based node-synchronization method. An alternative approach is a randomized algorithm based on the birthday paradox, which has been proposed for wireless sensor networks [34]. Specifically, the paper adopted the birthday paradox to develop a suite of protocols that save energy during the deployment phase of the sensors as well as the node discovery phase. The approach has been been adopted in a number of follow-up studies [49–51] including for a randomized data transfer protocol [35]. A comparative analysis of the two approaches (advertisement-based and randomized) is beyond the scope of this article.

9 CONCLUSIONS

In this article, we have presented a detailed design, implementation, and deployment of a wireless sensor network that collected river-temperature data at fine spatial and temporal scales. The overarching goal was to use temperature data, along with other data including fish habitat quality data, as well as hydrodynamic and geomorphic features to develop models of the forcings on river-temperatures. These models will better inform our understanding of juvenile salmonid energetics at ecological scales. In this article, we discussed the challenges of deploying the mesh network and how network storage was implemented to address intermittent link failures during inclement weather. We also described the design of the relay node and how a pseudo-thread architecture along with low power modes were leveraged to maximize battery life. We did not achieve 100% success, as a number of the sensors were destroyed during periods of heavy flow and by animals. Nevertheless, we were able to get considerable amounts of data, including particularly valuable and unique data from flood events. These data were valuable, because flood events are difficult to study, as they often scour and fill rivers and can render deployed logging sensors unrecoverable. As our data were uploaded in real-time, we obtained considerable data during these events. These data will help address questions about how river temperatures vary through both space and time and under differing flow regimes. This micro-habitat understanding of fluvial temperatures will inform future models of habitat quality for juvenile salmonids.

ACKNOWLEDGMENTS

The authors wish to express gratitude to the the entire lower Yuba River field crew from Pacific States Marine Fisheries Commission, without whom this study would have been impossible, notably Loren Stearman, Duane Massa, and John Cleveland. The extent to which Chelsea Martinez contributed to this work cannot be overstated; her impeccable organization and reliable field support were crucial to our success. We thank Richard Hambrick, who was instrumental in getting the development of this software off the ground and Megan Werdmuller von Elgg for her editorial feedback on this manuscript. Finally, we wish to thank the folks who helped us with this mountain of field work: Darcy Bostic, Annelise Del Rio, Arielle Gervais, Julius Henry, Sean Luis, Vanessa Lo, Seanna McLaughlin, Brianna Ordung, Lea Pollack, Kathy Russel, Rebecca Walker, Jason Wiener, Katie Woodworth, and Ken Zillig.

REFERENCES

[1] J. V. Ward. 1985. Thermal characteristics of running waters. In *Perspectives in Southern Hemisphere Limnology*. Springer Netherlands, Dordrecht, 31–46.

[2] B. W. Webb and F. Nobilis. 1997. Long-term perspective on the nature of the water temperature relationship: A case study. *Hydrol. Process.* 11 (1997), 137–147.

[3] K. Smith. 1975. Water temperature variations within a major river stream. *Hydrol. Process.* 6 (1975), 155–169.

[4] R. J. Danehy, C. G. Colson, K. B. Parrett, and S. D. Duke. 2005. Patterns and sources of thermal heterogeneity in small mountain streams within a forested setting. *Freshw. Biol.* 208 (2005), 287–302. DOI: http://dx.doi.org/doi:10.1016/j.foreco.2004.12.006

[5] N. J. Hetrick, M. A. Brusven, W. R. Meehan, and T. C. Bjornn. 1998. Changes in solar input, water temperature, periphyton accumulation, and allochthonous input and storage after canopy removal along two small salmon streams in southeast Alaska. *Trans. Amer. Fisher. Societ.* 127 (1998), 859–875. DOI: http://dx.doi.org/doi:10.1577/1548-8659(1998)127<0859:CISIWT>2.0.CO;2

[6] D. Caissie. 2001. The thermal regime of rivers: A review. *Freshw. Biol.* 51 (2001), 1389–1406. DOI: http://dx.doi.org/doi:10.1111/j.1365-2427.2006.01597.x

[7] P. A. Conrads and E. A. Roehle. 1999. Comparing physics-based and neural network models for simulating salinity, temperature and dissolved in a complex, tidally affected river basin. In *South Carolina Environmental Conference*.

[8] D. Caissie, N. El-Jabi, and G. Satish Mysore. 2001. Modelling of maximum daily water temperatures in a small stream using air temperatures. *J. Hydrol.* 251 (2001), 14–28. DOI: http://dx.doi.org/doi:10.1016/S0022-1694(01)00427-9

[9] M. B. Cardenas, J. W. Harvey, A. I. Packman, and D. T. Scott. 2008. Ground-based thermography of fluvial systems at low and high discharge reveals potential complex thermal heterogeneity driven by flow variation and bioroughness. *Hydrol. Process.* 22 (2008), 980–986. DOI: http://dx.doi.org/doi:10.1002/hyp.6932

[10] John S. Selker, Luc Thevenaz, Hendrik Huwald, Alfred Mallet, Wim Luxemburg, Nick van de Giesen, Martin Stejskal, Josef Zeman, Martijn Westhoff, and Marc B. Parlange. 2006. Distributed fiber-optic temperature sensing for hydrologic systems. *Wat. Resour. Res.* 42, W12202 (2006).

[11] Ian F. Akyildiz, Weilian Su, Yogesh Sankarasubramaniam, and Erdal Cayirci. 2002. Wireless sensor networks: A survey. *Comput. Netw.* 38, 4 (2002), 393–422.

[12] Jennifer Yick, Biswanath Mukherjee, and Dipak Ghosal. 2008. Wireless sensor network survey. *Comput. Netw.* 52, 12 (2008), 2292–2330.

[13] Mohamed Hefeeda and Majid Bagheri. 2009. Forest fire modeling and early detection using wireless sensor networks. *Ad Hoc & Sensor Wirel. Netw.* 7, 3-4 (2009), 169–224.

[14] Subramania Ananda Kumar and Paramasivam Ilango. 2018. The impact of wireless sensor network in the field of precision agriculture: A review. *Wirel. Person. Commun.* 98, 1 (2018), 685–698.

[15] Biljana L. Risteska Stojkoska and Kire V. Trivodaliev. 2017. A review of Internet of Things for smart home: Challenges and solutions. *J. Clean. Product.* 140 (2017), 1454–1464.

[16] Dae-Man Han and Jae-Hyun Lim. 2010. Smart home energy management system using IEEE 802.15. 4 and Zigbee. *IEEE Trans. Consum. Electron.* 56, 3 (2010), 1403–1410.

[17] J. Wang, D. Li, M. Zhou, and D. Ghosal. 2008. Data collection with multiple mobile actors in underwater sensor networks. In *28th International Conference on Distributed Computing Systems Workshops*. 216–221. DOI: http://dx.doi.org/10.1109/ICDCS.Workshops.2008.85

[18] Roland Kays, Margaret C. Crofoot, Walter Jetz, and Martin Wikelski. 2015. Terrestrial animal tracking as an eye on life and planet. *Science* 348, 6240 (2015), aaa2478.

[19] W. H. Chamberlain and H. L. Wells. 1879. *History of Yuba County, California, with Illustrations Descriptive of Its Scenery, Residences, Public Buildings, Fine Blocks and Manufactories*. Technical Report. Thompson & West.

[20] E. J. Eliason, T. D. Clark, M. J. Hague, L. M. Hanson, Z. S. Gallagher, K. M. Jeffries, M. K. Gale, D. A. Patterson, S. G. Hinch, and A. P. Farrell. 2011. Differences in thermal tolerance among Sockeye salmon populations. *Science* 332 (2011), 109–112. DOI: http://dx.doi.org/doi:10.1126/Science.1199158

[21] A. P. Farrell, N. A. Fangue, C. E. Verhille, D. E. Cocherell, and K. K. English. 2015. *Thermal Performance of Wild Juvenile Oncorhynchus Mykiss in the Lower Tuolumne River: A Case for Local Adjustment to High River Temperature*. Technical Report. Don Pedro Project.

[22] Antonio Cilfone, Luca Davoli, Laura Belli, and Gianluigi Ferrari. 2019. Wireless mesh networking: An IoT-oriented perspective survey on relevant technologies. *Fut. Internet* 11, 4 (2019). DOI: http://dx.doi.org/10.3390/fi11040099

[23] Guido R. Hiertz, Dee Denteneer, Sebastian Max, Rakesh Taori, Javier Cardona, Lars Berlemann, and Bernhard Walke. 2010. IEEE 802.11s: The WLAN mesh standard. *IEEE Wirel. Commun.* 17, 1 (2010), 104–111. DOI: http://dx.doi.org/10.1109/MWC.2010.5416357

[24] Renwei Huang, Zedong Nie, Changjiang Duan, Yuhang Liu, Liya Jia, and Lei Wang. 2015. Analysis and comparison of the IEEE 802.15.4 and 802.15.6 wireless standards based on MAC layer. In *Health Information Science*, Xiaoxia Yin, Kendall Ho, Daniel Zeng, Uwe Aickelin, Rui Zhou, and Hua Wang (Eds.). Springer International Publishing, Cham, 7–16.

[25] Martin Bor, John Edward Vidler, and Utz Roedig. 2016. LoRa for the internet of things. In *Proceedings of the 2016 International Conference on Embedded Wireless Systems and Networks*. 361–366.

[26] Pere Tuset-Peiró, Albert Anglés-Vazquez, José López-Vicario, and Xavier Vilajosana-Guillén. 2014. On the suitability of the 433 MHz band for M2M low-power wireless communications: Propagation aspects. *Trans. Emerg. Telecommun. Technol.* 25, 12 (2014), 1154–1168. DOI: http://dx.doi.org/10.1002/ett.2672

[27] Yordan Tabakov. 2017. Wizzilab: Connecting things/wizzimote. Retrieved from http://wizzilab.com/wiki/#!hardware/wizzimote.md#WizziMote.

[28] Texas Instruments. 2013. *CC430 Family User's Guide (Rev. E)*. Technical Report. Texas Instruments.

[29] Texas Instruments. 2018. *CC430F613x, CC430F612x, CC430F513x MSP430™ SoC with RF Core*. Technical Report. Texas Instruments.

[30] Maarten Weyn, Glenn Ergeerts, Luc Wante, Charles Vercauteren, and Peter Hellinckx. 2013. Survey of the DASH7 alliance protocol for 433 MHz wireless sensor communication. *Int. J. Distrib. Sens. Netw.* 9, 12 (Dec. 2013), 870430.

[31] Michael Buettner, Gary V. Yee, Eric Anderson, and Richard Han. 2006. X-MAC: A short preamble MAC protocol for duty-cycled wireless sensor networks. In *4th International Conference on Embedded Networked Sensor Systems*. 307–320.

[32] Tijs Van Dam and Koen Langendoen. 2003. An adaptive energy-efficient MAC protocol for wireless sensor networks. In *1st International Conference on Embedded Networked Sensor Systems*. 171–180.

[33] Tong Feng, Shilin Chen, Zhongke Feng, Chaoyong Shen, and Yi Tian. 2021. Effects of canopy and multi-epoch observations on single-point positioning errors of a GNSS in coniferous and broadleaved forests. *Rem. Sens.* 13, 12 (2021), 2325.

[34] Michael J. McGlynn and Steven A. Borbash. 2001. Birthday protocols for low energy deployment and flexible neighbor discovery in ad hoc wireless networks. In *Proceedings of the 2nd ACM International Symposium on Mobile Ad Hoc Networking & Computing*. ACM, 137–145.

[35] Jiahui Dai, Dmitry Degtyarev, Jingya Gao, Adrian Wang, Scott Burman, Ken Zillig, and Dipak Ghosal. 2020. Design and implementation of RAP-a randomized asynchronous protocol for data aggregation in wireless sensor networks. In *International Conference on Computing, Networking and Communications (ICNC)*. IEEE, 980–986.

[36] Joseph Polastre, Jason Hill, and David Culler. 2004. Versatile low power media access for wireless sensor networks. In *2nd International Conference on Embedded Networked Sensor Systems*. 95–107.

[37] Maarten Weyn, Glenn Ergeerts, Luc Wante, Charles Vercauteren, and Peter Hellinckx. 2013. Survey of the DASH7 alliance protocol for 433 MHz wireless sensor communication. *Int. J. Distrib. Sensor Netw.* 9 (2013). DOI: http://dx.doi.org/doi:10.1155/2013/870430

[38] Kemal Akkaya and Mohamed Younis. 2005. A survey on routing protocols for wireless sensor networks. *Ad Hoc Netw.* 3, 3 (2005), 325–349.

[39] Roger Alexander, Anders Brandt, J. P. Vasseur, Jonathan Hui, Kris Pister, Pascal Thubert, P. Levis, Rene Struik, Richard Kelsey, and Tim Winter. 2012. RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. DOI: 10.17487/RFC6550

[40] Dunkels Dunkels, Oliver Schmidt, Thiemo Voigt, and Muneeb Ali. 2006. Protothreads: Simplifying event-driven programming of memory-constrained embedded systems. In *4th ACM Conference on Embedded Networked Sensor Systems (SenSys'06)*.

[41] Alan Mainwaring, David Culler, Joseph Polastre, Robert Szewczyk, and John Anderson. 2002. Wireless sensor networks for habitat monitoring. In *1st ACM International Workshop on Wireless Sensor Networks and Applications*. 88–97.

[42] Yunhao Liu, Yuan He, Mo Li, Jiliang Wang, Kebin Liu, and Xiangyang Li. 2012. Does wireless sensor network scale? A measurement study on GreenOrbs. *IEEE Trans. Parallel Distrib. Syst.* 24, 10 (2012), 1983–1993.

[43] Xufei Mao, Xin Miao, Yuan He, Xiang-Yang Li, and Yunhao Liu. 2012. CitySee: Urban CO monitoring with sensors. In *IEEE INFOCOM Conference*. 1611–1619. DOI: http://dx.doi.org/10.1109/INFCOM.2012.6195530

[44] Holger Karl and Andreas Willig. 2007. *Protocols and Architectures for Wireless Sensor Networks*. John Wiley & Sons.

[45] Curt Schurgers, Vlasios Tsiatsis, Saurabh Ganeriwal, and Mani Srivastava. 2002. Optimizing sensor networks in the energy-latency-density design space. *IEEE Trans. Mob. Comput.* 1, 1 (2002), 70–80.

[46] Athanassios Boulis, Saurabh Ganeriwal, and Mani B. Srivastava. 2003. Aggregation in sensor networks: An energy–accuracy trade-off. *Ad Hoc Netw.* 1, 2-3 (2003), 317–331.

[47] Venkatesh Rajendran, Katia Obraczka, and Jose Joaquin Garcia-Luna-Aceves. 2006. Energy-efficient, collision-free medium access control for wireless sensor networks. *Wirel. Netw.* 12, 1 (2006), 63–78.

[48] Arvind Kumar et al. 2017. *Energy Efficient Clustering Algorithm for Wireless Sensor Network*. Ph.D. Dissertation. Lovely Professional University.
[49] Santashil PalChaudhuri and David B. Johnson. 2002. Birthday paradox for energy conservation in sensor networks. *Sleep* 9 (2002), 14mA.
[50] Qiang Wang, Andreas Terzis, and Alex Szalay. 2010. A novel soil measuring wireless sensor network. In *IEEE Instrumentation & Measurement Technology Conference*. IEEE, 412–415.
[51] Tingjun Chen, Javad Ghaderi, Dan Rubenstein, and Gil Zussman. 2018. Maximizing broadcast throughput under ultra-low-power constraints. *IEEE/ACM Trans. Netw.* 26, 2 (2018), 779–792.

Received 27 February 2021; revised 20 May 2022; accepted 24 May 2022

Identifying dominant effects on sensed river temperatures in a regulated, dry-summer subtropical river

Abstract

River water temperatures can generally be characterized to warm as they move downstream. Beyond this simple trend, river water temperatures can be difficult to characterize, particularly at fine spatial and temporal scales (few meters and diel or less, respectively). Channels are warmed by insolation, which in turn is mediated by vegetative shading and velocities. Hyporheic flows are less susceptible to warming. Yet, with climate change raising river temperatures, decreasing precipitation, and enhancing drought, dry-summer subtropical rivers are especially vulnerable and require further scrutiny of their spatial and temporal temperature dynamics. To build on existing literature, we addressed three novel research questions: (1) At which temporal scales did fluvial temperatures vary most in a regulated, dry-summer subtropical river? (2) Do hyporheically connected side-channels experience warming, or do their main-channel connections limit these effects? (3) Over short reaches of relatively cold water, is the downstream warming trend still prevalent? In an effort to address these questions, we designed and implemented a novel wireless monitoring network in the lower Yuba River (LYR) in California's Central Valley, and then installed 35 temperature sensors in the river to collect temperature data every 15 minutes in late 2018 through mid-2019. In order to ascertain the dominant effects driving the sensed river temperatures, we created several Bayesian models, which divided the river geographically, but were not

spatially explicit. We found that models that incorporated the interactive effect of a diel and day-of-sampling effects with random walk priors yielded the best fit. Within the study reach, our model results indicated that temporal effects dominated the temperature data, with diel-effects less important than longer, multi-day effects. We found that even with strong hyporheic connections, the sampled side channel was warmer than the main-channel. We also found only minimal downstream warming within our 3 km study reach. Given how velocities differed between upstream and downstream sensors as well as between mainstem and side-channel sensors in the study reach, we infer a strong velocity effect best addressed by spatially explicit models. Given the strong observed temporal effects, we anticipate the need to model interactions in spatio-temporal models of river temperatures.

Background and Significance

Fluvial temperature

Dynamics and patterns in fluvial temperatures are complex: forced by a number of different river and environmental processes. These processes in-turn interact in ways that vary across both spatial and temporal scales. The broadest general trend in fluvial temperature is well characterized: both mean and variance of temperatures increase from headwaters to mouth (Ward 1985). Beyond the broad generalization of downstream warming, fluvial temperatures are difficult to measure and model at spatial and temporal scales relevant to biota. Generally, insolation warms the river during the day (Danehy et al 2005, Daniels and Danner 2020), creating a diel return-phase signal, but this insolation can be mediated locally by shading from riparian vegetation (Johnson

2004, Jackson et al 2021). Seasonal inputs from precipitation and snowmelt cool the channel (Danehy et al 2005). Hyporheic flows – non-surface flows moving through river substrate – moderate temperatures in small areas of the channel (Ebersole et al 2003), but are mediated by topography and morphology (Bilby 1984). Further, sediment can modulate water temperature through localized conduction, which in turn depends upon the sizes and composition of channel sediments (Story et al 2003). Each of these processes interacts with each other locally, and at different spatial and temporal scales, and on managed rivers, these effects are further mediated by dam management (Daniels and Danner 2020). These processes are difficult to observe, and instead, their complex interactions yield observable and measurable river temperature.

Most studies of water temperature prior to the 1980s used a Six's thermometer, which measures maximum and minimum temperature at the point of deployment between researcher resets of the apparatus. The observed warming from headwaters downstream to the river mouth (Macan 1958, Ward 1985, Danehy et al 2005, Daniels and Danner 2020) can mean different things based upon sampling methodology and reset-frequency of a Six's thermometer. For example, an increase in average temperature could have been the result of either a higher maximum or minimum temperature. These changes in observed extrema say nothing about changes in the underlying distributions from which they are sampled. Rivers could have identical extrema, but wildly different skews and variances, and one could imagine a temperature regime that remained at or near either extreme for much of the sampling period – this would be a very different river from one that remained near the average and only

reached these extrema for short durations. Further, the scale of measurement changes the meaning of these readings if the thermometer is reset daily, seasonally, or annually.

The general downstream warming trend varies temporally, and is confounded by a variety of factors including stream size, bed slope, and timescale of observation. In small mountain (Danehy et al 2005) or flat (Macan 1958) streams, at daily timescales, downstream sites had higher maximum temperatures than upstream sites, with little or no downstream change in minimum temperatures. In a study of a river in northern England, downstream sites were found to have higher mean monthly and annual temperatures. In this system, maximum monthly temperatures were relatively similar upstream and downstream, but minimum monthly temperatures were considerably higher downstream (Smith 1975). This difference (minimum temperature was more variable than was maximum temperature) may have been the result of the size of this study system, as larger rivers have been found to be less impacted by air temperatures (Danehy et al 2005), and the elevation change of the river in question, as Johnson (1971) found a strong negative relationship between elevation and temperature. The general trend of fluvial temperatures increasing while moving downstream is impacted by seasonal maxima, with the strongest trends observed at summer thermal maxima (Hopkins 1971, Smith 1975). The trend has been found to be best observed at hourly or daily timescales (Smith 1975). These findings were generalized by Ward (1985), who found that maximum temperatures increased from headwaters to mouth, with annual ranges also increasing downstream (Ward 1985).

The temporal scales of fluvial temperature are further complicated by the relationship between air and water temperatures (Daniels and Danner 2020). If this relationship were particularly strong, it would follow that diel water temperatures would vary most in summer, when air temperatures vary the most over the course of a day, and to vary least in the winter, when air temperatures vary least over the course of the day, particularly in the mild winters of Mediterranean climates (Andrew and Sauquet 2017). Yet, in small streams, one study found that diel variability in fluvial temperature was greatest in spring and least in autumn (Macan 1958). Elsewhere, as expected, the greatest diel variability in water temperature was found in summer, and the least variable in winter (Hopkins 1971). Regardless, the relationship between air and water temperatures was found to be correlative and not causative (Webb and Nobilis 1997). The strength of the relationship between air and water temperatures varied with the timescale over which it was investigated: monthly mean temperatures were found to be strongly correlated with air temperatures, but annual mean temperatures were not (Webb and Nobilis 1997). These results further complicate any generalized understanding of fluvial water temperatures. The variable relationship between air and water temperatures could be the result of the seasonality of intensity of solar insolation. In other words, solar radiation warms both water and air, but the air does not warm moving waters via convection.

The impact of solar insolation on river temperature is high, but the effect is heterogeneous and localized (Jackson et al 2021). The effect of insolation is greatest when waters are slower moving and shallower, which increases the duration of exposure and could explain the downstream warming trends described above (Daniels

and Danner 2020). This solar warming was found to be mitigated by increased flow, localized vegetative shading, and snow-melt (Smith 1975, Beschta and Taylor 1988, Daniels and Danner 2020, Jackson et al 2021).

Conditions in rivers in Mediterranean climates are particularly variable (Andrew and Sauquet 2017, Cid et al 2017, Daniels and Danner 2020), and are likely to be particularly susceptible to perturbations from climate change (Andrew and Sauquet 2017). This was thought to be the result of the nature of Mediterranean climates, which experience most of their precipitation during winters and with almost no summer precipitation (Cid et al 2017). Which is to say that precipitation-based cooling aligns temporally with seasonal radiative, conductive, and convective effects, exacerbating them rather than mitigating them (as summer-wet climates might). Seasonal and Interannual temperature effects are difficult to predict in Mediterranean rivers as the duration and magnitude of precipitation are more variable than in other humid climates (Andrew and Sauquet 2017, Cid et al 2017).

Study purpose

Given the critical gaps in the existing scientific literature regarding temporal variability in water temperature in different parts of a river as described in the review provided above, the overall scientific purpose of this study was to investigate fluvial water temperature spatial and temporal complexity in a regulated, perennially flowing, lowland river in the dry-summer subtropical climate using novel data collection methods and Bayesian time-series modeling. This overall purpose was turned into the following three specific, tractable research questions:

1. (RQ1) At which temporal scales, ranging from hours to months, do fluvial temperatures vary most in a regulated, dry-summer subtropical river?

2. (RQ2) Do temperatures in hyporheically connected side-channels differ in temperature from the main channels?

3. (RQ3) Over short (3 km) reaches of relatively cold water, fed by releases from deep-water dams upstream, is the downstream warming trend still prevalent, or is it best observed over longer river reaches?

The first research question RQ1 was our primary question, from which the sampling and models were designed. The other two research questions RQ2 and RQ3 were answered using results that came out of the main inquiry.

Contributions

We found strong, interacting temporal forcings on fluvial temperatures. These effects dominated other modeled effects. We saw no meaningful downstream warming effect over this short study-reach, and observed that the side-channel was only marginally warmer than the mainstem despite being largely disconnected. We attributed this to the strong hyporheic connections between mainstem and side channel.

Study Area

The lower Yuba River (LYR) is a multi-threaded river in California's Central Valley. The river is constrained by engineered training-berms, has a cobble bed, and is predominantly fed by snow-melt, keeping the river colder than many other rivers in

California's Central Valley. The LYR merges with the Feather River near Marysville, CA. The Feather in turn flows downstream to the Sacramento River, which flows to the San Francisco Bay (Figure 1).

The study area was a ~3 km reach immediately downstream from Daguerre Point Dam, an ~8 m tall dam, built in the early 1900s to divert water for nearby agriculture. One consequence of this dam is that the water is well-mixed in a plunge pool that in turn feeds the reach just below the dam (Figure 1). This reach was channelized and constrained by manmade training-berms until a series of floods in 2017 and 2018 scoured out a secondary channel that went around one of these training berms, reconnecting with the main-channel about 1 km downstream. The side channel below this merge point remained hyporheically connected, but was not connected by surface flows except during flood events (Figure 1). Just below the study reach is the Marysville hydrologic monitoring station. In 2010-2012, temperature sensors in the fish ladders at Daguerre Point Dam, just above the study region averaged between 10 and 15 ^{o}C for each month in those three years (Figure 2).

Topo-bathymetric digital elevation models of the LYR corridor were previously produced at ~ 1-m resolution in 2008, 2014, and 2017 (Gervasi et al 2021). Given these topo-bathymetric models, meter-resolution, two-dimensional hydrodynamic models were made for the whole LYR for each time as well using the methods of Pasternack (2011), enabling an understanding of the study area's hydraulic processes (e.g. Abu-Aly et al 2013; Strom et al 2016; Barker et al 2018). In-channel riverbed morphological units

were subsequently objectively mapped for each of those using the method of Wyrick et al (2012).

Methods

Field methods

In order to study fluvial temperatures within the study reach, we deployed a wireless sensor network consisting of 35 temperature sensors located in both the primary and side-channels both above and below the merge-point. The wireless sensor network, and its design, deployment, and data collection are thoroughly explained in Burman et al (2020) and Burman and Gao et al (2023), but the most pertinent details relevant here are as follows. Sensors were attached (five each) to sensor nodes, which contained all of the necessary electronics and wireless hardware. Water temperature was measured about four times per hour (sensor nodes did not have real-time clocks so sampling frequency was temperature dependent, fluctuating between ~14 and 16 minutes). The present paper covers the temporally explicit models generated from this data to answer the posed research questions.

It was imperative to measure temperatures at different locations within the wetted channel as well as at a suitable frequency in time. Temperatures needed to be assayed along the length of the study reach as well as across the channel. In wide gravel and cobble rivers with diverse morphological units and lateral groundwater discharges from hillsides, river temperatures can vary more across a channel than longitudinally.

To address *RQ1*, we measured temperatures multiple times per hour. Each sensor node sampled its associated five sensors every ~ 15 minutes, yielding roughly four samples per sensor, per hour. Further, given the strong seasonal forcings associated with river temperatures and the seasonal variabilities in river flows, we endeavored to sample for at least a year to capture the totality of seasonal fluctuations, which would have allowed us to identify a full seasonal cycle of temperatures. Unfortunately, a river enhancement project (Southall et al 2022) brought our sampling to an end prematurely – we instead consider the hourly, daily, and monthly temporal signals (RQ1).

To address *RQ2*, we sampled the mainstem, side-channels connected at upstream and downstream ends, and other off-channel habitats (Figure 1). Finally, to address *RQ3* regarding downstream warming trends, we sampled longitudinally along the length of the study reach.

Sensor locations were selected to sample temperatures in a diversity of river morphologies, while balancing worker safety. In order to sample a wide diversity of hydrologic conditions, areas of the river were selected in advance of field trips on the computer using the extensive LYR geo-database and considering their geographic location and geomorphology. Once in the field, workers selected exact sites and transect orientations where sensors could safely be installed as planned or as near to the previously selected locations as possible given real-world constraints. Whenever practicable, a random number generator was used to set the number of steps taken from the edge of the study region to begin the sensor installation.

Among the seven sensor nodes that survived the majority of the sampling season, five (1-4, 6) were in the northern perennial side channel and two (5, 7) were in the mainstem baseflow channel (Figure 1) as follows:

- Sensor node 1 traversed a deep pool with clear water at the downstream end of the side channel. The pool contained aquatic vegetation where depths were suitably shallow. For discharges below ~ 12,000 cfs, the lower half of the side channel was entirely hyporheically fed. Also, the downstream end of the deep pool was well connected to the mainstem channel, so it was perpetually wetted. The pool was ~ 50 m across. The sensors installed here were spaced ~ 10 m apart. This sensor survived for the duration of the study.

- Sensor node 2 was in a riffle just upstream of the deep pool. The sensors on this sensor string were spaced 5 m apart. The furthest upstream sensor on the string was placed in an extremely cold, but very small stream that was generally disconnected from the rest of the side channel. It was entirely hyporheically fed. This tiny channel then fed into a riffle in the larger side channel. The second to last sensor was in the deepest part of the pool in the riffle-pool couplet, and the final sensor was on a shallow bar. This sensor survived for the duration of the study.

- Sensor node 3 was in a glide within the side channel. Near the northern channel margin there was a small pool and at the other end (in the deepest part) was a slightly faster glide. Sensors were spaced 5 m apart. This sensor survived for the duration of the study.

- Sensor node 4 was in a very small riffle that at summer baseflow was very slow moving, could easily be walked by workers, and was often just a few centimeters in depth. Just upstream from this small riffle was a large pool. Sensors were spaced about 5 m apart. This sensor node survived for the duration of the study.

- Sensor node 5 was placed in a fast glide in the main channel. This area was wide, flat, and fast moving with very loose coarse sediment: sensors were spaced 10 m apart. It was extremely difficult for workers to hold position while wading, and as such, we could not collect precise GPS coordinates at each sensor point, and instead interpolated from the end points of the sensor string knowing their spacings. This sensor was destroyed by a flood event on 3 March 2019.

- Sensor node 6 was in an extremely high velocity region of the river, which was difficult for workers to traverse. Its sensors were spaced 5 m apart. This upper side channel was directly fed by the main channel, immediately below Daguerre Point Dam. This area had previously been a solid gravel mass in the year before the deployment, but a major flood in 2018 scoured out the gravel and reconnected this area of the river (Gervasi et al 2021). This sensor extended from an off-channel pool, out into the side channel, over a small riffle, but in what was mostly a faster-moving glide. This sensor survived the majority of our sampling, but was destroyed on 23 February 2019.

- Sensor node 7 was installed at the merging of two channels, one smaller and intermittently wetted (often little more than a few cm in water depth), the other

was the very high velocity main channel. During installation, workers struggled in this area with the high velocities. Sensors on this node were spaced 10 m apart. This also coincided with a meander bend. Here, the sensor traversed a fast moving riffle and glide, and then on the northern side, it curved into a slower moving pool. This sensor appeared to have been destroyed by a bear on 7 March 2019.

Sensors were polled by the attached nodes every 15 minutes, and data traversed the network and were uploaded to a server at the University of California, Davis campus. Sensors were damaged and disabled throughout this period, with some remaining online only until February 2019, while others remained through May 2019, when a river restoration project began, which required the network to be taken down (Burman and Gao et al 2023).

Analytical methods

In order to explain the factors that force fluvial temperatures, we sought to generate statistical models of varying complexity. By comparing these models, we can ascertain which factors best explained the observed temperatures. Before making any inference about effects on river temperature, we needed to identify the model that best fit the sensed river temperature data. This identification allowed us to infer which factors best explained the sensed temperatures. Once selected, we used the best-fitting model to ascertain posterior temperature distributions for each factor.

We fit eight linear Bayesian models, with increasing levels of complexity, then assessed the tradeoff between model fit and model complexity. In this way, the best-fitting model and its underlying priors gave us insight into the river system sampled. By building Bayesian models, we could specify priors for temporal effects that incorporated the autocorrelation inherent in the temporal data. For each temporal effect, we tested different random walk priors, with either a first-order or a second-order lag. The amount of lag defined the amount of smoothing in the prior, with higher order priors characterizing smoother temporal effects, and therefore greater temporal autocorrelation. The greater flexibility of Bayesian models allowed for more robust interactions between model effects. Once we identified the best fitting model, which yielded insight into our first research question (RQ1), we considered the posterior distributions for each effect. Notably, models were stratified to find the effect of sensor position in the main stem and side channel (RQ2) as well as upstream and downstream (RQ3).

We attempted to fit eight different models, beginning with simple models, with increasing numbers of interactions and increasingly complex temporal effects. We also tried to improve model fit by experimenting with different priors for each of the temporal effects within the models (Table 1).

Temperature models

Temperature models were fit using Integrated Nested Laplace Approximation (explained further in the next section below; Rue et al 2009, Martins et al 2013). The selected model split the wireless sensor network into four quadrants, separated by their position

upstream or downstream of the merge-point in our two-mile river reach, and whether the sensor was in the side-channel or the main-channel (Figure 1). Following the established literature (Rue et al 2009, Martins et al 2013, Gómez-Rubio 2020), the temporal effects were included as random walk priors, with the order determined experimentally using a sensitivity analysis.

These eight models were:

1. Model 1: This is a hierarchical model with sensor number (1-5) nested within sensor node (1-7). This model was structured such that there were no shared effects between sensor numbers (eg: 1_1, 1_2,...8_5). That is to say, sensor 1 within sensor node 1 was distinct from sensor 1 on sensor node 2 and so on. This structured the individual sensors within their respective nodes. We hypothesized that the main and side channels would be better differentiated than were the up- and downstream regions of the river. As such, we treated the upstream/downstream effect as a grouping effect (Gaussian Kronecker Markov Random Field: GKMRF; explained in the next section) within the main and side channel effect. We also treated both of these as random effects, with an independent and identically distributed prior. This created a per-region effect for the upstream and downstream as well as the main and side channel. These hierarchical, IID effects were retained for all subsequent models.

2. Model 2: The second model extended the simpler nested model and incorporated a day-of-year term as a cyclic random walk prior. We defined the sampling period

within a 365-day cycle, in the hopes that doing so would allow some ability to model the remainder of the year without directly sampling it.

3. Model 3: The third model was very similar to the day-of-year model (2), but instead of using the day of the year, we used the hour of the day as a cyclic random walk prior.

4. Model 4: This model combined the effects (linearly – not the interaction between them) of the hourly effect and the 365-day day-of-year effect, both with cyclic random walk prior.

5. Model 5: In this model, we retained the hourly effect, modeled a cyclic random walk prior. We replaced the 365-day day-of-year effect with a simpler, day-of-sampling effect modeled with a random walk prior. Given that our sampling spanned less than a full year, we considered that the return-phase of a year was perhaps meaningless.

6. Model 6: In this model, we removed the hourly effect and looked at just the day-of-sampling effect with a random walk prior.

7. Model 7: In this model we parameterized the day-of-sampling effect with a random walk prior, and created a GKMRF grouping variable with the cyclic hour-of-day effect, parameterized by a cyclic random walk prior. This created a grouping variable for hour of day within day of year, which is in turn an interaction effect between hour of day and day-of-sampling period.

8. Model 8: In this model we adjusted model 7, altered by parameterizing the hour-of-the-day effect as a scaled variable instead of cyclic variable. This was due to INLA not allowing for scaling of cyclic GKMRF grouping variables. We selected the best fitting model (methods below) and all further analysis was done using that model. The posterior distribution for each effect from this model were then sampled and used in further analyses.

Integrated Nested Laplace Approximation

Integrated Nested Laplace Approximation (INLA) is a method for fitting Bayesian models using approximation of Latent Gaussian Models. INLA is distributed as an R-package (R-INLA; Rue et al 2009, Martins et al 2013) that allows for the fitting of most Bayesian models. Instead of using Monte Carlo Markov Chains (MCMC) as used in BUGS, JAGS, and STAN, INLA instead uses integration to solve complex Bayesian models. Because there are no large Markov chains, INLA is considerably faster than MCMC methods. Further, INLA uses the PARDISO library (Alappat et al 2020, Bollhöfer et al 2019, Bollhöfer et al 2020). This library is a thread-safe solver for systems of linear equations. Because it is thread safe, INLA is inherently parallel, conferring the further advantage of parallelism to solve models even faster.

INLA has a number of unique features, which were implemented in these models. It allows for the specification of priors that include independent, identically distributed effects (IID) as well as spatially and temporally stratified effects, which can be structured as autoregressive effects (AR: with a defined lag), or as random walks with lag of one or two (RW1, RW2; Rue et al 2009, Martins et al 2013). All temporal effects were modeled

as random walk effects, as these priors have been well established as good priors for temporal effects, and we established the order of these random walks experimentally (Zuur et al 2017, Gomez-Rubio 2020).

INLA allows for complex interactive effects by defining groupings between effects. This feature generates a second model for between-group correlation (Gomez-Rubio 2020). The result is a variance structure that is the Kronecker product of the individual effects. This Gaussian Kronecker Markov Random Field (GKMRF) reflects the inherent dependence of the effects (Riebler 2011). Such product random fields are conceptually similar to interaction effects, but are preferable in INLA (Riebler 2011).

We experimented with a number of different hyper-parameters and priors for all terms in the model. After extensive testing, we selected those in Zuur et al (2017) for the random walk effect's prior (θ, which is the log-precision) equal to the the Penalized Complexity precision prior (PC: equations 1 and 2, from Rue et al 2009, Martins et al 2013), with hyper-parameters specified as u=5 (or ~3x the sample standard deviation for the response variable (~1.635)) and α=0.01.

$$\pi(\tau) = \frac{\lambda}{2}\tau^{-3/2} * e^{-\lambda\tau^{-1/2}}, \ \tau > 0 \quad (1)$$

for λ > 0, where

$$\lambda = -\frac{ln(\alpha)}{\mu} \quad (2)$$

IID effects were specified with priors using the same PC prior and hyperparameters, however, the prior for the IID effects was specified as the precision, and not the log-precision per (Zuur et al 2017, Gomez-Rubio 2020).

Validation methods

We sought to validate the models against the observed data. Often, this is done with a training set and a validation set of some subset of the observed data. We opted to instead use INLA's two internal validation methods, which provide a validation metric that is computed while the model is fit. These methods are well studied and characterized, so we opted not to do an additional validation step.

1. *Conditional predictive ordinate* (CPO) is a cross-validation criterion for model assessment. For each observation, the CPO is the posterior probability of observing that observation when the model is fit using all data but that point. Conceptually this is very similar to a leave-one-out validation for all fitted points. Higher values indicate better fit for that point (Petit 1990; Held, Schrödle, and Rue 2010; Gómez-Rubio 2020).
2. *Predictive integral transform* (PIT) measure, which is the probability, for each observation, of a new value being lower than the actual observed value. Well fitting models will have a distribution close to a uniform distribution between 0 and 1 (Marshall and Spiegelhalter 2003; Held, Schrödle, and Rue 2010; Gomez-Rubio 2020).

Model selection criteria

Deviance Information Criterion (DIC) is conceptually similar to the ubiquitous Akaki's Information Criterion, and considers fit, while penalizing model complexity via effective number of parameters. Lower values of DIC are preferred (Spiegelhalter et al.2002, Gomez-Rubio 2020). The Watanabe-Akaike Information Criterion (WAIC) also known as widely applicable Bayesian Information Criterion differs from DIC only in how the number of effective parameters are computed (Watanabe 2013; Gelman, Hwang, and Vehtari 2014). In this study we used the DIC and WAIC model selection criteria to determine which model fit the data best, then all further analysis was done using that model. The posterior distribution for each effect was plotted and used to address the three questions.

Testing the Research Questions

The testing of the research questions was multi-faceted. We considered eight different candidate models summarized in Table 1. Invariably, the models which best fit the data, both absolutely (validation methods: CPO and PIT) and penalized for model complexity (DIC and WAIC), will suggest the strongest forcings of fluvial temperature, which addresses RQ1. That is to say that if the dominant temporal effect on fluvial temperatures in the study reach is time-of-day or day-of-sampling (or year), then we would expect the corresponding model to yield the best fit, and for the relevant increase in complexity to be worthwhile in comparative model heuristics (DIC and WAIC). Further, if these effects interact, which would yield a significant increase in effective parameters (~sampling days * 24hrs/day), the corresponding increase in explanatory

power, would need to be huge to offset this. Alternatively, if temporal effects have only minimal effect on fluvial temperatures, we would expect the model one to have the best fit, particularly when penalizing complexity as it is estimated to have only a few dozen effective parameters (the 35 sensors are stratified, with some overlap with the geographic stratification).

Once these models were fit, we considered the credible intervals of the posterior temperatures from the best-fitting model to address the secondary questions and to quantify the effects of main and side channel sensor placement (RQ2) as well as up and downstream (RQ3). Non-overlapping credible intervals can reasonably be interpreted to be sampled from different distributions.

Results

Model heuristics and selection

We found that our simplest model, which considered the nested sensors within the sensor nodes, and were then placed into the GKMRF of main and side-channel within up and downstream effects had the poorest fit, despite its simplicity. This was true even when considering model selection criteria that penalize model complexity (Model 1; DIC & WAIC ~585,000; Table 1). It is important to remember that DIC and WAIC assess model fit, and they are penalized for model complexity. This model can be thought of as a baseline, with each subsequent comparison amounting to a cost-benefit analysis for further model complexity.

Adding to the nested model the hour of the day improved the fit (Model 3; DIC & WAIC ~573,000; Table 1), but adding in the DoS (day of sampling) or DoY (day of year) improved the model fit far more (Models 2&6; DIC & WAIC ~335,000; Table 1). Notably, there were far more sampled days than there were hours in the day – despite this complexity, we still found the day of the sampling to be a better predictor of temperature, whose complexity was offset by improved fit. We did notice that using the day of sampling yielded a slightly better fit than did day-of-year. This was likely because day of sampling was a simpler metric, and benefits of modeling a cyclic return-phase (like day of year) is contingent upon having multiple complete cycles sampled.

Models including both day-of-sampling (or year) and hour-of-the-day fit better still (Models 4&5; DIC & WAIC ~277,000; Table 1). But, models that used the GKMRF of the hour-of-the-day within the day-of-sampling, fit best (Models 7&8; DIC & WAIC ~250,000; Table 1). The complexity of these models (~40-100 times more effective parameters than in the simplest, nested model (Model 1); Table 1) was offset by their improved fit.

Temperature models with the day of year (or day of sampling) as an effect were best modeled by a first order random walk (rw1) prior (Table 1). The temperature models containing hour of day as an effect were less sensitive to order of the prior for this effect (Table 1). In the more complex GKMRF (where the effective number of parameters is the product of parameters from the day of sampling and hour of day) the additional smoothing of the second order random walk yielded a considerable reduction in the effective number of parameters – making these models preferable (Table 1).

We found Model 7, which was the GKMRF of the cyclic hour of the day (rw2) and the day of sampling (rw1) fit best (Table 1). This model's complexity was the result of complex temporal interactions, that these interactions improved fit – despite penalizing complexity – yields insight into our primary research question regarding the temporal variability in river temperatures. In further considering the fit of Model 7, the residuals appeared almost entirely independent (Figure 3), as they followed a horizontal line of slope zero (red line). The observed and expected values (Figure 4) had an excellent, one-to-one relationship (red line of slope one: any appearance otherwise is the result of truncated axes to ensure that detail is visible). Broken out into uniform quantiles, with normal quantiles, the observed PIT scores fit well to a straight line of slope one (Figure 5, black line has slope one), though PIT residuals did deviate from the expected uniform distribution (0:1), and instead were center-biased (Figure 6). The distribution of the residuals of CPO scores were, as expected, skewed toward one (Figure 7). As Model 7 yielded the best fit, we derive all further analyses from just this model.

Modeled temperature – spatial and temporal effects

We saw little variation among all sensor nodes, indicating that the majority of the observed variability in water temperature was explained elsewhere – and was not the result of unexplained within-node or between-node variation (Figure 8). Within-node spreads had consistent 95% credible intervals of approximately 0.4°C (Figure 8), indicating low within-node variation. Across all sensors, the spread of the mean posterior temperature was less than 0.2°C (Figure 8), suggesting that the inter-node effects were well explained by the other model components.

The effect of sensor placement in the main stem or side channel showed a posterior mean of the main channel being about 1.5 ^{o}C cooler than the side-channel (Figure 9). The credible intervals of the posterior distribution also do not overlap (Figure 9). The grouped effect of upstream and downstream showed minimal effect, with the downstream sensors slightly warmer (< 0.5^{o}C) than the upstream sensors (Figure 10).

The primary effect we observed was temporal. The greatest variability in water temperature was the result of the GMKRF of the time of day and the day of the sampling period (Figure 11). We also found consistently low standard deviations in temporal effect on fluvial temperatures (Figure 11).

Discussion

Given the considerable increase in model complexity of the chosen model, relative to the simplest, hierarchical model considered (~40-100 times the effective number of parameters), the inclusion of temporal effects was imperative. In particular, these temporal effects dominated the posterior temperature, particularly when considering their interaction. River temperatures should be considered within the context of both seasonal and diel effects (Table 1, Figure 11). Our model choice underscores the important interactions between these effects.

Moving beyond the simple model selection, we considered three questions about fluvial temperatures in the lower Yuba River, a regulated, gravel and cobble river in a dry-summer subtropical climate in California's Central Valley (Figure 1). We consider

our secondary questions first, before returning to a more thorough consideration of our primary question regarding dominant temporal effects.

Prevalence of downstream warming trend (RQ3)

We observed a minimal downstream warming trend (< 0.5°C) within our stratification, between our upstream and downstream sensor locations (Figure 10). This may have been because of the short (~ 3 km) length of the study reach, which was only about 60 times as long as the main channel was wide (~ 50 m). Notably, the upstream and downstream sites were quite different depending upon where in the reach they were. Put simply, the downstream side channel was the outlier relative to the remainder of the sampling locations. It was not connected by surface flows to the main channel except during flood events. The fact that this unique sampling area was relegated to the downstream side channel only likely played a role in obfuscating a downstream warming trend. Anecdotally, workers observed that when the main stem meandered closer to the training berm and the side-channel, the side channel tended to be colder, with higher velocities. Sensor two was at such a site (Figure 1). This suggests a significant hyporheic connection under the training berm.

Temperature difference in the hyporheically connected side-channels (RQ2)

The side channel was warmer than the main stem, the difference between mean sampled posterior temperatures was approximately 1.5°C relative to the main channel (Figure 9). Without the hyporheic connection between the main stem and side channels, we would have anticipated an even larger difference in temperature given the orders of

magnitude difference in volume and velocity of water flowing in the main stem and side channels. The main stem was about four to six times the width of the downstream side channel, and about double the width of the upstream side channel.

In this study, we chose to model two dominant strata within our sampling: (a) upstream versus downstream, and (b) side channel versus main stem. We then created an interactive effect between these strata for our models – yielding in effect four distinct groups (ie: upstream side channel, upstream main stem, downstream side channel, downstream main stem). As is always the case in modeling studies, there is a trade-off between characterizing the local system and generalizing to address broader trends. By stratifying as we did, we inherently opted for the more general. Had we instead created a single stratified effect between the lower velocity downstream side channel relative to the remainder of the sensors, we may have been able to better characterize the temperature differences between them, and may well have found more explicit results about how temperatures varied between them. We may have better characterized the downstream warming within the main channel. But, such a model would have yielded a result specific only to this system, which is why we did not do that. Instead, the approach taken has more value as a novel scientific contribution that future studies can use in other settings.

Dominant temporal scales of temperature variability (RQ1)

Temperature does not vary just seasonally, or just daily. Rather, these fluctuations in temperature interact (Figure 11). Temporal effects dominated in our selected model of fluvial temperature. As would be expected, we saw extrema in diel temperatures at

around 2 am and 2 pm (Figure 11). We observed a minimum temperature in February and a maximum in November and May (the warmest months sampled; Figure 11). That said, while we see considerable variation in temperatures through both day of sampling and hour of day, we note that incorporating sampling day improved the model fit more than did adding sampling hour (using all information criteria and validation methods; Table 1). Considering their interaction yielded our best model fit (Table 1). We observe consistent and low sample standard deviations, indicating strong predictive trends throughout the sample period (Figure 11). We also observe none of the seasonal or diel fluctuations in observed variability described in past work (eg: Macan 1958, Hopkins 1971).

Past work has considered temporal scales over which temperature changes are best observed in rivers, and whether changes in mean temperature were the result of changes in one or both extrema over a particular sampling period. These were particularly important questions when water temperatures were predominantly sampled using Six's thermometers (eg: Macan 1958, Hopkins 1971). Work that has focused on identifying trends in temperature variability over singular temporal durations (eg: daily, monthly, or seasonal) scales likely missed the interactions between these timescales, which we find characterize much of the true variability in fluvial temperature. To this end, while past work using Six's thermometers identified changes in the maximum and minimum temperature observed at a focal timescale, we find no such pattern, instead observing seasonal trends in both diel maxima and minima over the course of our sampling period.

We also note that there were days in which diel fluctuations were lost (Figure 11, eg: 13 Feb 2019). These days were uniformly cold: there was minimal warming over the course of the day. This underscores the fallibility of samplings that fail to incorporate both daily and seasonal fluctuations. The date of 13 February 2019 was a day with persistent rain. It was overcast throughout the day (National Weather Service Forecast Office 2022). On high-precipitation, overcast days such as this, the river is cooled directly by the precipitation and insolation is limited by cloud cover. In locations like California's Central Valley, overcast rainy conditions are typically limited to winter months, which can contribute complexity to the interactions between daily and seasonal cycles.

We observed an obvious seasonal signature within our temperature data, but all of these effects were smaller than those of singularly cold days, which in fact have no meaningful daily maximum temperature, and instead remain cold all day (Figure 11). We interpret that these signatures, with anomalously cold days suggest that the temperatures within this reach were dominated by precipitation and cloud cover, which is a seasonal phenomenon but also driven by subseasonal weather cycles, interannual ocean-atmosphere climate cycles, and are indistinguishable-from-random occurrences.

While we were unable to sample for a complete year, we see no reason to expect changes to these patterns in the summer months. We instead find that in warmer seasons, fluvial temperatures were forced higher throughout the day, and in cooler seasons, fluvial temperatures were forced lower throughout the day. While we found the standard deviation in posterior temperatures from our models to be relatively consistent

throughout our sampling period, there are exceptions. Notably, near the end of our sampling period we observed numerous overcast days, which caused the batteries to die in our solar-powered data-aggregation point. This data loss decreased our sampling density, and increased the model variance (Figure 11).

Our models do not support any general or consistent pattern in fluvial temperature variance as most of the temporal regions with increased variance were the result of periods of poor sampling due to power losses during inclement weather. However, in the first week of March 2019, we observe notable increases in sampled temperature deviations (Figure 11). This coincided with the end of a flood, during which the flow was roughly five times (CDEC 2022: 709 m^3/s on 26 February 2019) the bankfull flow (142 m^3/s). During the early days of March 2019,when the flow was receding to the bankfull flow, this pulse of cold water at the end of February, with falling flows likely altered the local dynamics, and in particular, impacted different sensors differently – with the sensors away from the hyporheically fed side channel likely experiencing much more extreme temperature differences. This time period also was characterized by considerable cloud-cover (National Weather Service Forecast Office 2022). Without insolation, the falling flows from this flood event likely dominated in forcing temperatures for main-channel sensors, hence the considerable increase in variability in this time period. Spatial models would provide considerable insight into these dynamics.

Study Outcomes Significance

This study serves to extend the understanding of the dominant temporal scales over which temperature fluctuations in a regular, dry-summer subtropical river in California's

Central Valley. Rivers in this type of climate are subject to greater temperature variability than rivers in other climates. Identifying the dominant temporal scales of these variations will be crucial in modeling changes in these systems, particularly as they are expected to be acutely impacted by climate change. Results should be relevant to regulated rivers in the Mediterranean region, western south America, South Africa, and south and western Australia.

Future Study

We expect that spatially explicit models, which consider proximity to shore and level of vegetative shading would yield better insights into the exact dynamics at play. Indeed, if cloud cover (or conversely, insolation) dominates the temperature effect within this reach, then the cooling effect on overcast days should be greatest in areas of the river with slower flows, less hyporheic flow, and less vegetative shading, and should have the least effect in fast-moving areas, with hyporheic flows, and in areas with the most vegetative cover. Our future work will expand upon these models to incorporate the spatial structure of the sampled reach.

Acknowledgements

This work was supported by a cooperative agreement (F16AC00735) with the Anadromous Fisheries Restoration Program within the US Fish and Wildlife Service with the support of Dr. Beth Campbell and Paul Cadrett of the USFW. The lower Yuba River field crew from the Pacific States Marine Fisheries Commission were crucial in

completing the relevant fieldwork, notably Loren Stearman, Duane Massa, and John Cleveland. Chelsea Martinez's contribution to this work cannot be overstated, this project would have failed without her diligent work, impeccable organization, and reliable field support in a challenging working environment. Richard Hambrick was instrumental in getting the development of network software off the ground. This work was taken over by Jingya Gao, a brilliant undergraduate student in computer science who worked independently and collaboratively to bring the networking software to fruition. Megan Werdmuller von Elgg provided crucial editorial support in the drafting of this document. The folks who helped with the mountain of field work for this study, bringing it to fruition were Darcy Bostic, Annelise Del Rio, Arielle Gervais, Julius Henry, Vanessa Lo, Sean Luis, Seanna McLaughlin, Brianna Ordung, Lea Pollack, Kathy Russel, Rebecca Walker, Jason Wiener, Katie Woodworth, and Ken Zillig.

Literature Cited

Abu-Aly TR, Pasternack GB, Wyrick JR, Barker R, Massa D, Johnson T. 2013. Effects of LiDAR-derived, spatially-distributed vegetative roughness on 2D hydraulics in a gravel-cobble river at flows of 0.2 to 20 times bankfull. Geomorphology. doi:10.1016/j.geomorph.2013.10.017.

Alappat C, Hager G, Schenk O, Thies J, Basermann A, Bischop A, Fehske H, Wellein G. 2020. A Recursive Algebraic Coloring Technique for Hardware-Efficient Symmetric Sparse Matrix-Vector Multiplication, ACM Transactions on Parallel Computing, Vol. 7, No. 3(19)

Andrew JT, Sauquet E. 2017. Climate Change Impacts and Water Management Adaptation in Two Mediterranean-Climate Watersheds: Learning from the Durance and Sacramento Rivers. Water.9.126: doi:10.3390/w9020126

Barker, J., Pasternack, G. B., Bratovich, P., Duane, M., Wyrick, J. R., Johnson, T. 2018. Kayak drifter surface velocity observation for 2D hydraulic model validation. River Research and Applications. DOI: 10.1002/rra.3238.

Beschta RL, Taylor RL. 1988. Stream temperature increases and land use in a forested Oregon watershed. J Amer Water Resour As. 24:19–25. doi:10.1111/j.1752-1688.1988.tb00875.x.

Bilby RE. 1984. Characteristics and frequency of cool-water areas in a western Washington stream. J Freshwater Ecol. 2:593–602. doi:10.1080/02705060.1984.9664642.

Bollhöfer M, Eftekhar Ai, Scheidegger S, Schenk O. 2019 Large-Scale Sparse Inverse Covariance Matrix Estimation, SIAM J. Sci. Comput., 41(1), A380–A401

Bollhöfer M, Schenk O, Janalik R, Hamm S, Gullapalli K. 2020. State-of-The-Art Sparse Direct Solvers, Parallel Algorithms in Computational Science and Engineering. Modeling and Simulation in Science, Engineering and Technology, 1-32, Birkhäuser

Burman SG, Gao J, Ghosal D, Pasternack G, Fangue N. 2020. Design, Implementation, and Deployment of TempMesh: A Wireless Mesh Network to Aggregate River-Temperature Data. 2020 IEEE International Conference on Sensing, Communication and Networking (SECON Workshops). 10.1109/SECONWorkshops50264.2020.9149773.

Burman SG, Gao J, Pasternack GB, Fangue NA, Cadrett P, Campbell E, Ghosal D. 2023. TempMesh – A Flexible Wireless Sensor Network for Monitoring River Temperatures. ACM Trans Sens Netw. https://doi.org/10.1145/3542697

California Data Exchange Center (CDEC). 2022. Yuba River Near Marysville. California Department of Water Resources. Accessed 18 Jul 2022. [https://cdec.water.ca.gov/dynamicapp/QueryF?s=MRY&d=04-Mar-2019+00:00&span=144hours]

Cid N, Bonada N, Carlson SM, Grantham TE, Gasith A, Resh VH. 2017. High Variability Is a Defining Component of Mediterranean-Climate Rivers and Their Biota. Water. 9(52).doi.org/10.3390/w9010052

Danehy RJ, Colson CG, Parrett KB, Duke SD. 2005. Patterns and sources of thermal heterogeneity in small mountain streams within a forested setting. Forest Ecol Manage. 208:287–302. doi:10.1016/j.foreco.2004.12.006.

Daniels ME, Danner EM. 2020. The drivers of river temperatures below a large dam. Water Resour Res. 56. doi.org/10.1029/2019WR026751

Ebersole JL, Liss WJ, Frissell CA. 2003. Cold water patches in warm streams: physicochemical characteristics and the influence of shading. J Amer Water Resour As. 39:355–368. doi:10.1111/j.1752-1688.2003.tb04390.x.

Gelman A, Hwang J, Vehtari A. 2014. Understanding Predictive Information Criteria for Bayesian Models. Stat Comput. 24 (6): 997–1016. https://doi.org/10.1007/s11222-013-9416-2.

Gervasi AA, Pasternack GB, East AE. 2021. Flooding duration and volume more important than peak discharge in explaining 18 years of gravel-cobble river change. Earth Surface Processes and Landforms 46 (15): 3194-3212. DOI: 10.1002/esp.5230

Gómez-Rubio V. 2020. Bayesian Inference with INLA. Chapman & Hall/CRC Press. Boca Raton, FL.

Held L, Schrödle B, Rue H. 2010. Posterior and Cross-Validatory Predictive Checks: A Comparison of MCMC and INLA. In Statistical Modelling and Regression Structures – Festschrift in Honour of Ludwig Fahrmeir, edited by T. Kneib and G. Tutz, 91–110. Berlin: Springer Verlag

Hopkins CL. 1971. The annual temperature regime of a small stream in New Zealand. Hydrobiologia. 37:397–408. doi:10.1007/BF00018811.

Jackson FL, Hannah DM, Ouellet V, Malcolm IA. 2021. A deterministic river temperature model to prioritize management of riparian woodlands to reduce summer maximum river temperatures. Hydrol Processes. doi.org/10.1002/hyp.14314

Johnson FA. 1971. Stream temperatures in an alpine area. J Hydrol.14:322-336

Johnson SL. 2004. Factors influencing stream temperatures in small streams: substrate effects and a shading experiment. Can J Fish Aquat Sci. 61:913–923. doi:10.1139/f04-040.

Macan TT. 1958. The temperature of a small stony stream. Hydrobiologia. 12:89–106. doi:10.1007/BF00034143.

Marshall EC, Spiegelhalter DJ. 2003. Approximate Cross-Validatory Predictive Checks in Disease Mapping Models. Stat Medi. 22(10): 1649–1660.

Martins TG, Simpson D, Lindgren F, Rue H. 2013. Bayesian computing with INLA: New features. Comput Stats Data An. 67:68-83. doi:10.1016/j.csda.2013.04.014.

National Weather Service Forecast Office. 2022. Marysville 2019, Sacramento, CA. Accessed 18 Jul 2022. [https://www.wrh.noaa.gov/climate/temp_graphs.php?wfo=sto&p=precipitation&year=2019&stn=KMYV]

Pasternack, G. B. 2011. 2D Modeling and Ecohydraulic Analysis. Createspace: Seattle, WA.

Pettit LI. 1990. The Conditional Predictive Ordinate for the Normal Distribution. J R Stat Soc. Ser B (Methodological) 52(1): 175–184.

Riebler A. 2011. Gaussian Kronecker product Markov random fields. University of Zurich, Division of Biostatistics

Rue H, Martino S, Chopin N. 2009. Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with discussion). J Royal Stat Soc, Ser B. 71(2):319-392

Southall N, Sawyer A, Diaz S, Hammersmark C, Bowles C. 2022. Golden opportunity. Civil Eng. 92 (2): 50-59. https://ascelibrary.org/doi/epdf/10.1061/ciegag.0001610

Smith K. 1975. Water temperature variations within a major river stream. Nord Hydrol. 6:155-169

Spiegelhalter DJ, Best NG, Carlin BP, Van der Linde A. 2002. Bayesian Measures of Model Complexity and Fit (with Discussion). J R Stat Soc, Ser B. 64(4):583–616.

Story A, Moore RD, Macdonald JS. 2003. Stream temperatures in two shaded reaches below cutblocks and logging roads: downstream cooling linked to subsurface hydrology. Can J Forest Res. 33:1383–1396. doi:10.1139/x03-087

Strom, M. A., Pasternack, G. B., Wyrick, J. R. 2016. Reenvisioning velocity reversal as a diversity of hydraulic patch behaviors. Hydrologic Processes, 30 (13): 2348-2365. doi: 10.1002/hyp.10797.

Ward JV. 1985. Thermal characteristics of running waters. In: Davies BR, Walmsley RD, editors. Perspectives in Southern Hemisphere Limnology. Dordrecht: Springer Netherlands. p.31–46

Watanabe S. 2013. A Widely Applicable Bayesian Information Criterion. J Mach Learn Res. 14: 867–897.

Webb BW, Nobilis F. 1997. Long-term perspective on the nature of the water temperature relationship: a case study. Hydrol Process. 11:137-147

Wyrick JR, Pasternack GB. 2012. Landforms of the Lower Yuba River. Prepared for Yuba Accord River Management Team. University of California, Davis, Davis, CA, 91pp. DOI: 10.13140/RG.2.2.27716.63365

Zuur AF, Ieno EN, Saveliev AA. 2017. Beginner's Guide to Spatial, Temporal and Spatial-Temporal Ecological Data Analysis with R-INLA. Highland Statistics Ltd. Newburgh, UK

Tables

Table 1. A comparison of statistics regarding fit and selection for models and priors attempted in this study (DoY/DoS: Day of Year/Day of Season, rw: random walk prior)

	Priors		DIC			WAIC			
Model	DoY/ DoS	Hour	DIC	Saturated	# Params	WAIC	# Params	CPO	PIT
Sensors nested within nodes with up/down stream and main/side channel interacting (iid) effects	-	-	585155	3766188	32	585150	27	0.1911	0.4884
Sensors nested within nodes with up/down stream and main/side channel interacting (iid) effects, plus day of year	rw1	-	335990	3517029	217	333990	211	0.4334	0.4991
	rw2	-	336007	3517047	213	334001	203	0.4333	0.4991
Sensors nested within nodes with up/down stream and main/side channel interacting (iid) effects, plus hour of day	-	rw1	572940	3753974	58	572931	49	0.1987	0.4867
	-	rw2	572938	3753972	57	572930	49	0.1987	0.4867
Sensors nested within nodes with up/down stream and main/side channel interacting (iid) effects, plus day of year and hour of day	rw1	rw1	278984	3460023	241	274699	225	0.5272	0.5021
	rw1	rw2	278984	3460024	239	274761	219	0.5264	0.5021
	rw2	rw1	279000	3460040	237	274738	225	0.5269	0.5021
	rw2	rw2	279005	3460044	233	274735	219	0.5269	0.5021
Sensors nested within nodes with up/down stream and main/side channel interacting (iid) effects, plus day of year and hour of day	rw1	rw1	278768	3459807	239	274456	223	0.5279	0.5014
	rw1	rw2	278767	3459807	239	274393	217	0.5285	0.5014
	rw2	rw1	278778	3459817	237	274472	217	0.5276	0.5014
	rw2	rw2	278776	3459815	236	274471	216	0.5277	0.5014
Sensors nested within nodes with up/down stream and main/side channel interacting (iid) effects, plus day of sampling	rw1	-	333958	3514997	217	331826	203	0.4372	0.5005
	rw2	-	333971	3515010	215	331856	209	0.4370	0.5005
Sensors nested within nodes with up/down stream and main/side channel interacting (iid) effects, plus interaction between day of sampling and hour of day	rw1	rw1	250139	3431179	2014	245280	2451	0.5883	0.5014
	rw1	rw2	250084	3431124	1492	244367	1497	0.5872	0.5014
	rw2	rw1	250595	3431635	2008	246278	2655	0.5865	0.5014
	rw2	rw2	250538	3431578	1572	244796	1586	0.5871	0.5014
Sensors nested within nodes with up/down stream and main/side channel interacting (iid) effects, plus day of year and hour of day, but with non-cyclic hourly effect	rw1	rw1	250209	3431249	1988	245323	2357	0.5869	0.5014
	rw1	rw2	250381	3431420	1392	244462	1301	0.5864	0.5014
	rw2	rw1	250712	3431752	2004	246331	2596	0.5859	0.5014
	rw2	rw2	250999	3432039	1484	245159	1420	0.5854	0.5014

Figures

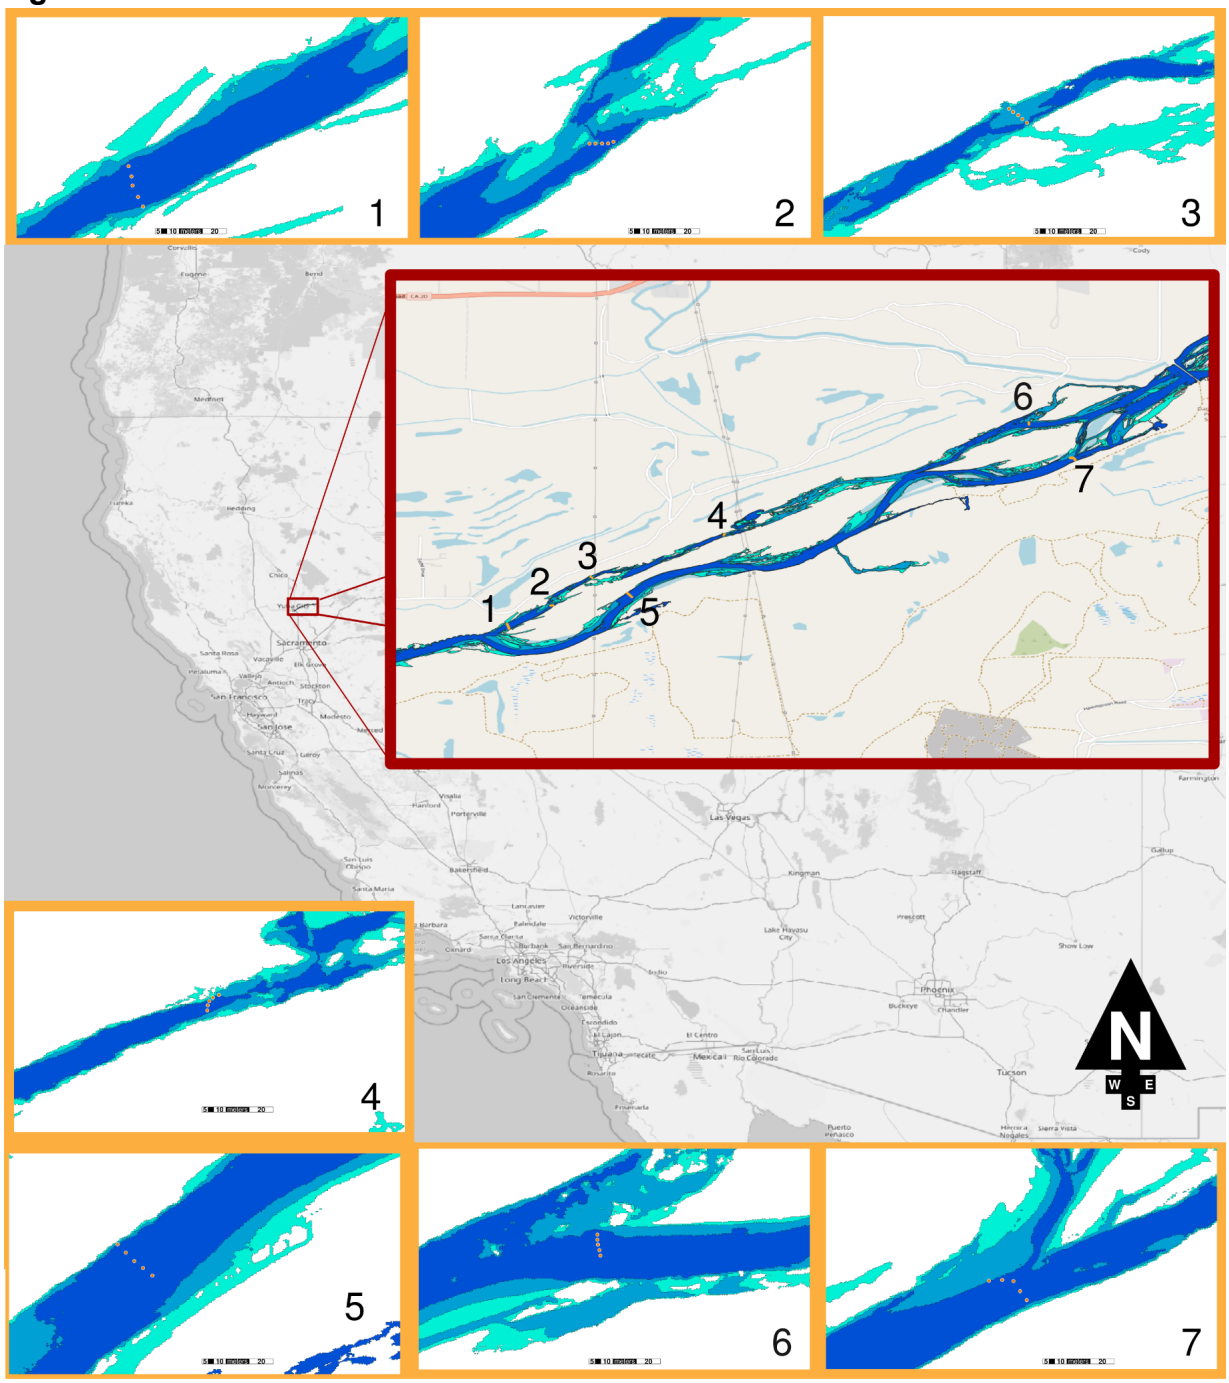

Figure 1. Map of the study reach of the lower Yuba River. In the top-right of the inset map, is Daguerre Point Dam. The river flows from the top right to the bottom left. Orange lines correspond to sensors (1-7), which span the channel width at that point. Red dots represent the exact location of the individual sensors, measured on GPS, except for sensor 4, which was subject to flows that were too deep and fast for precise GPS measurement, and points were therefore estimated based upon end point locations.

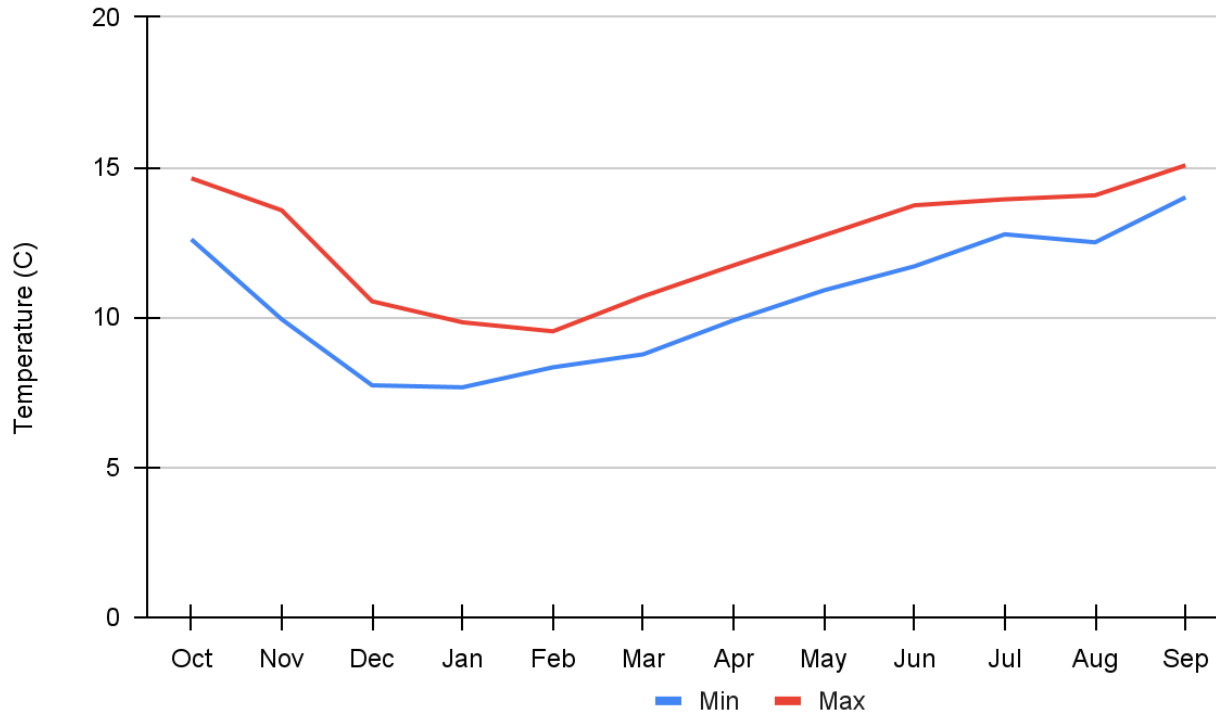

Figure 2. Maximum and minimum temperatures, by month, averaged across 2010, 2011, and 2012, adapted from Yuba County Water Agency 2017.

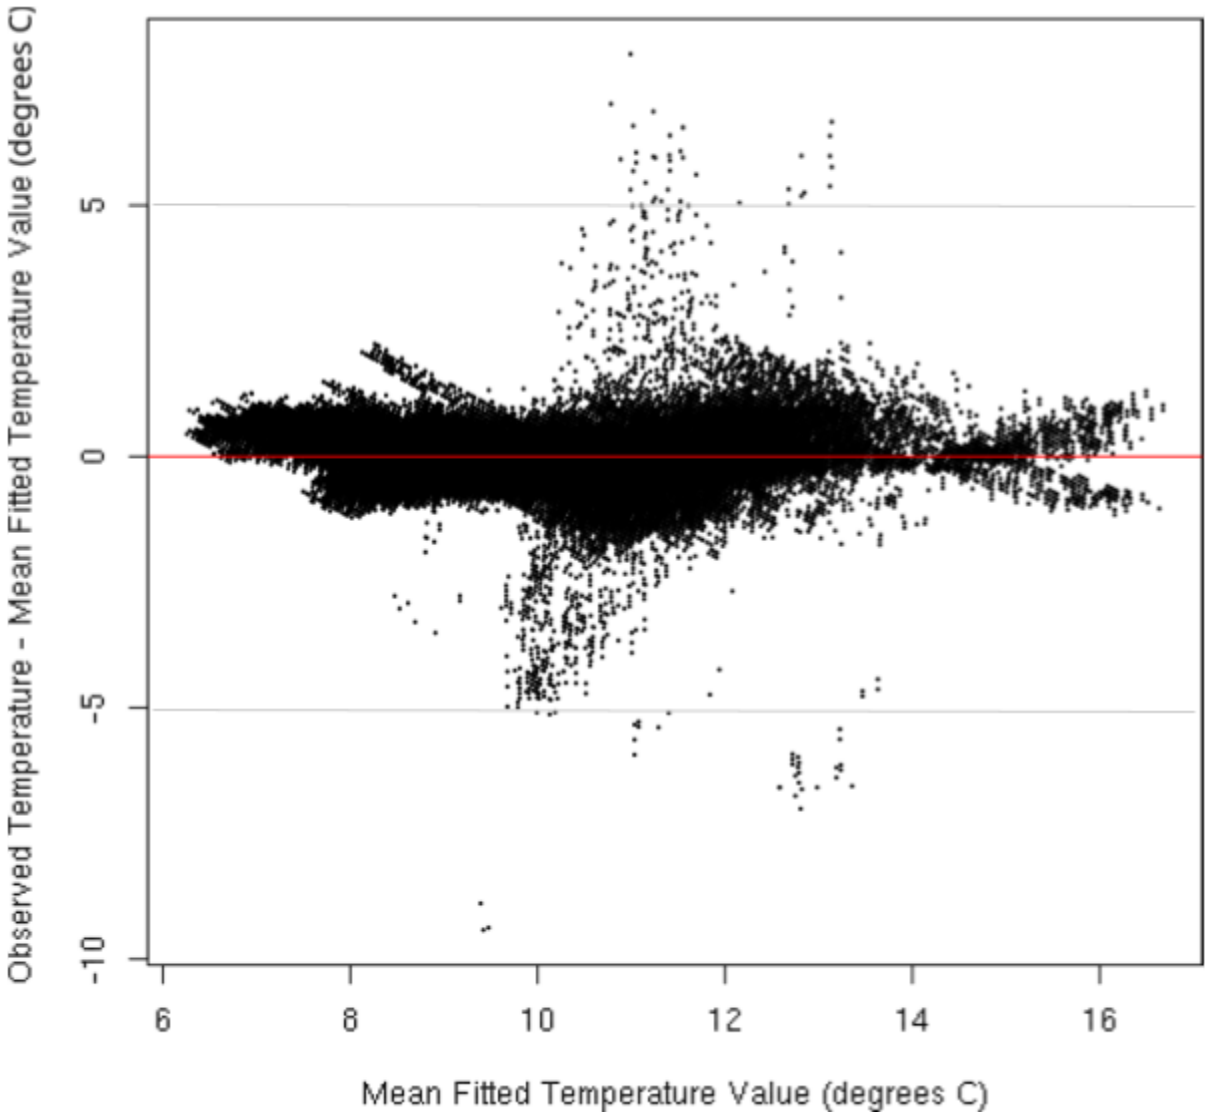

Figure 3. Residual temperatures (observed-fitted) relative to fitted temperature value. Were this point-cloud to fit a line with any slope, the result would indicate dependence of model deviation with respect to mean fitted temperature. Ideally, this point cloud would fit tightly to a line of slope zero (red line).

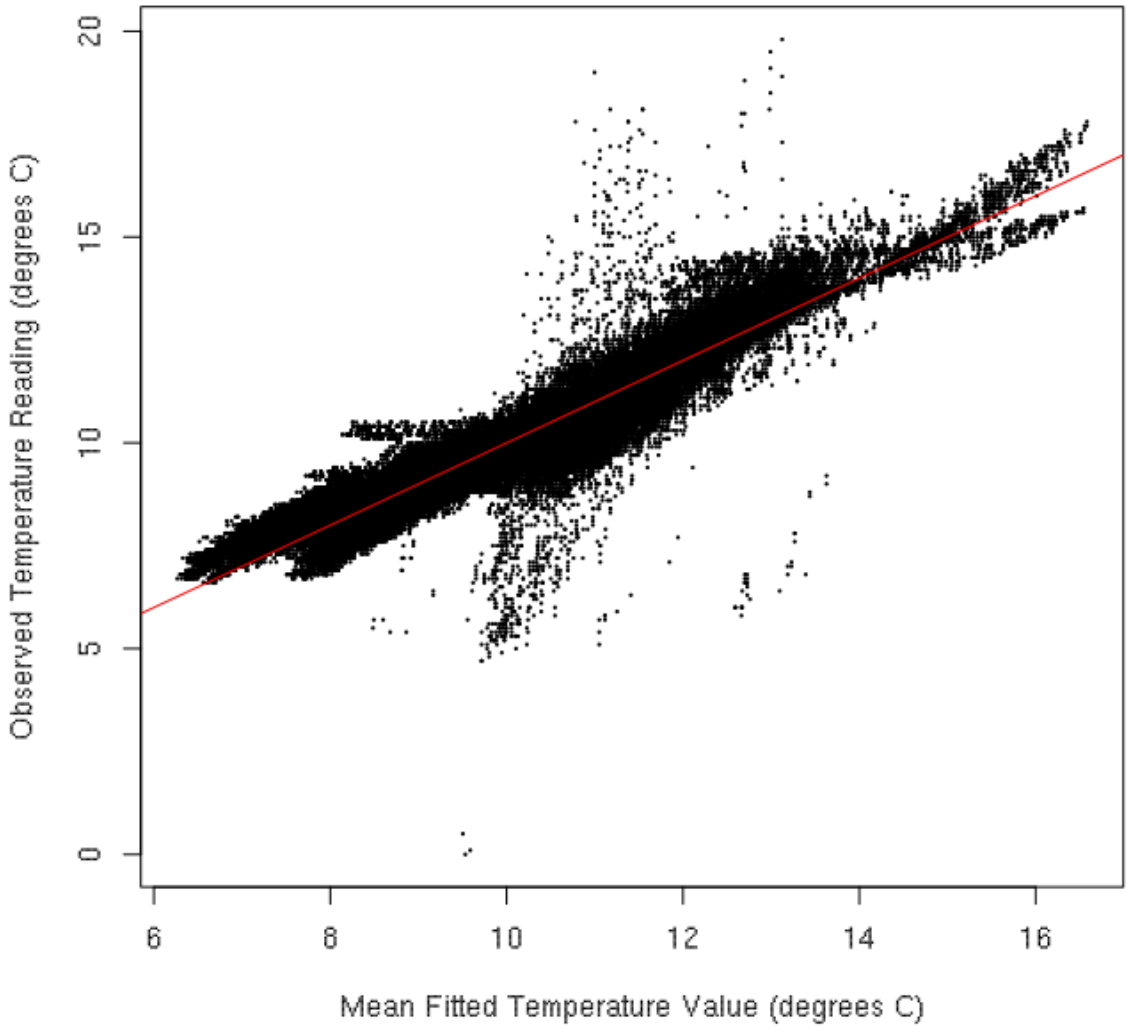

Figure 4. Observed temperature reading relative to mean fitted temperature. The better the model is fit, the closer the point cluster should fall to a straight line of slope one (red line). The red line may appear askew, this is only because axes were truncated to show more detail within the point cloud.

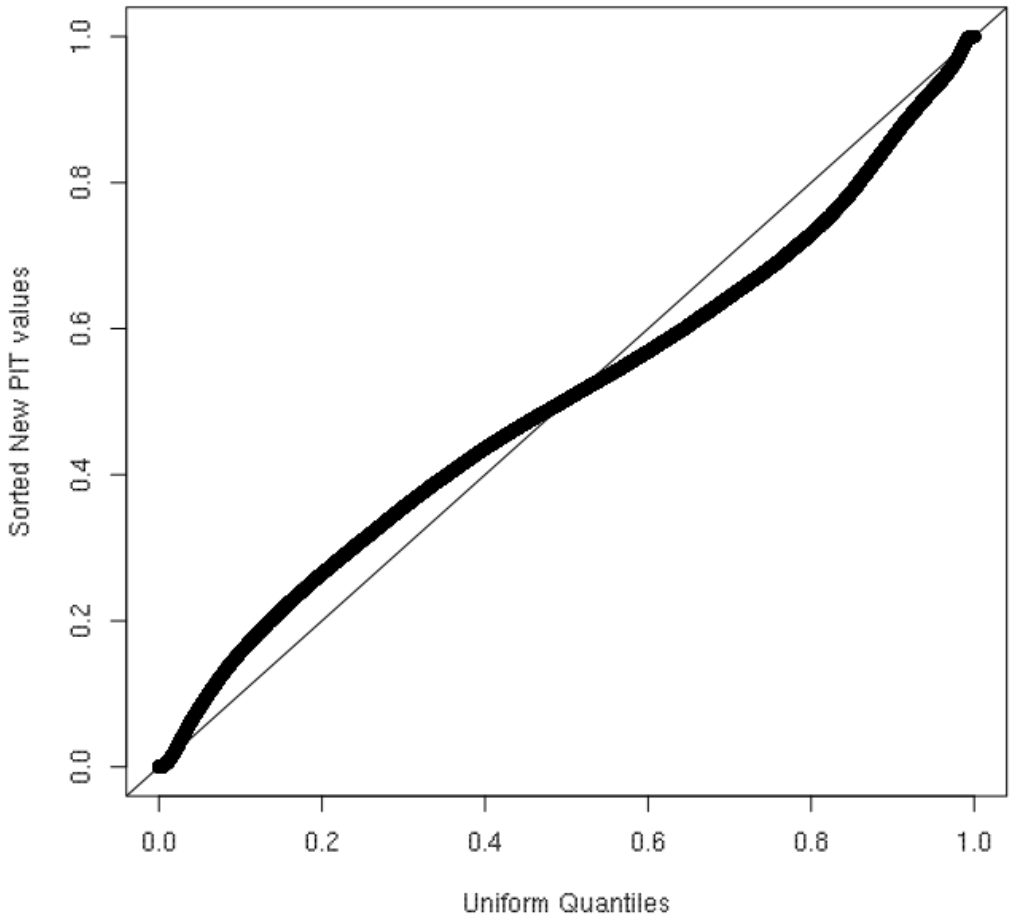

Figure 5. A comparison of the uniform quantiles of PIT scores, relative to sorted new PIT values. This should, ideally follow a straight line of slope one.

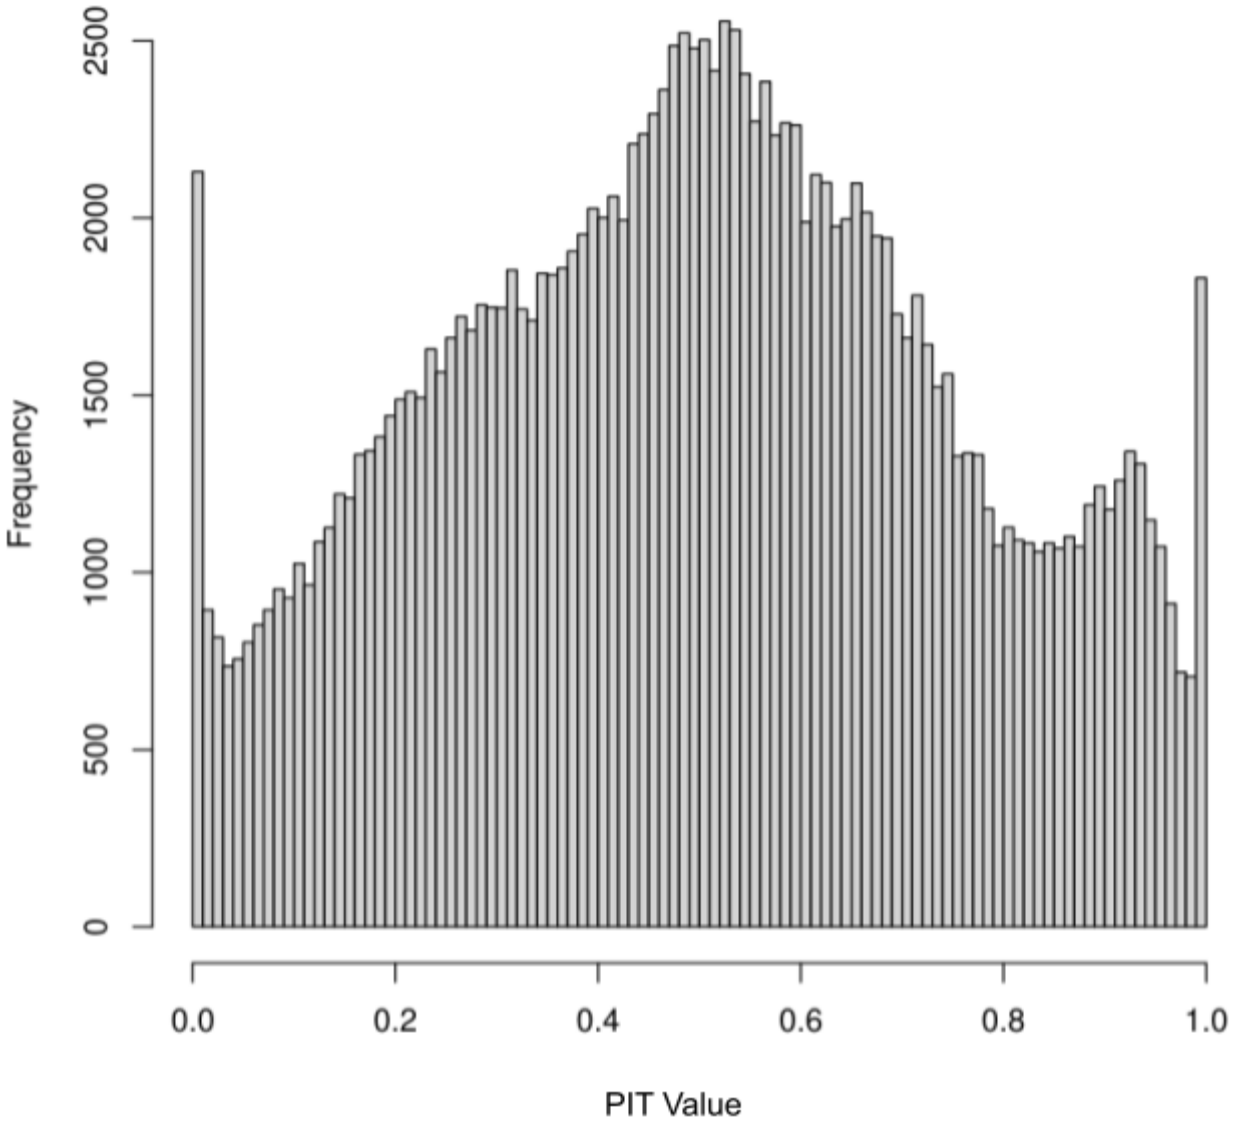

Figure 6. Frequency of “PIT values”, which are the probability of obtaining a lower value than the observed value in the data, when leaving out that value from the model fit. Ideally, the result would be a uniform distribution between 0 and 1. Crucially, the mean should also be very close to 0.5.

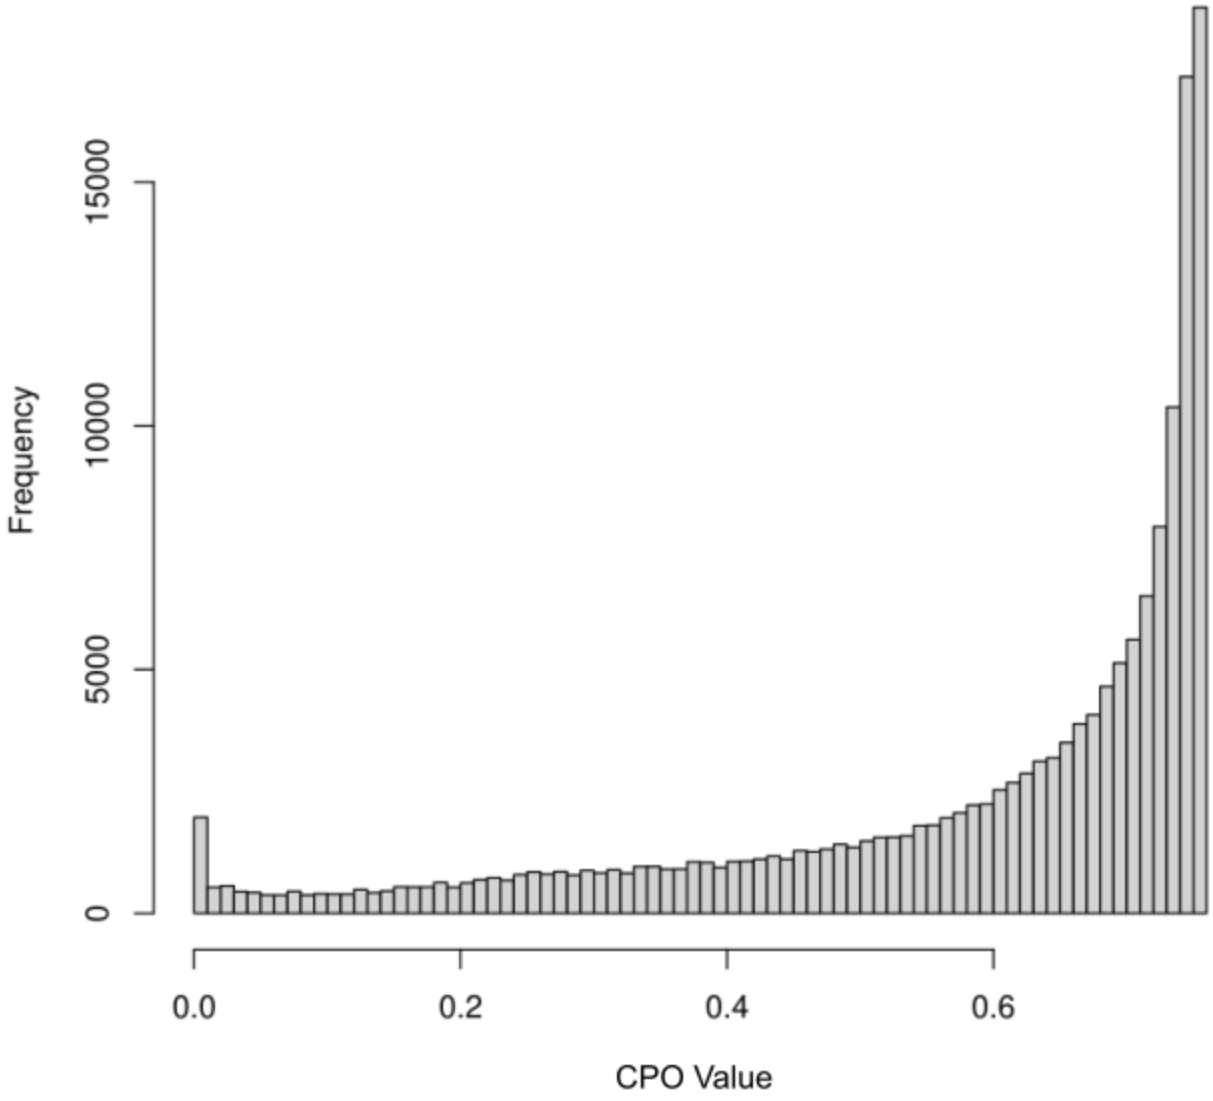

Figure 7. Frequency of “CPO values”, which are the posterior probability of observing that observation when the model is fit using all data but that point. These should be skewed towards one in well-fit models.

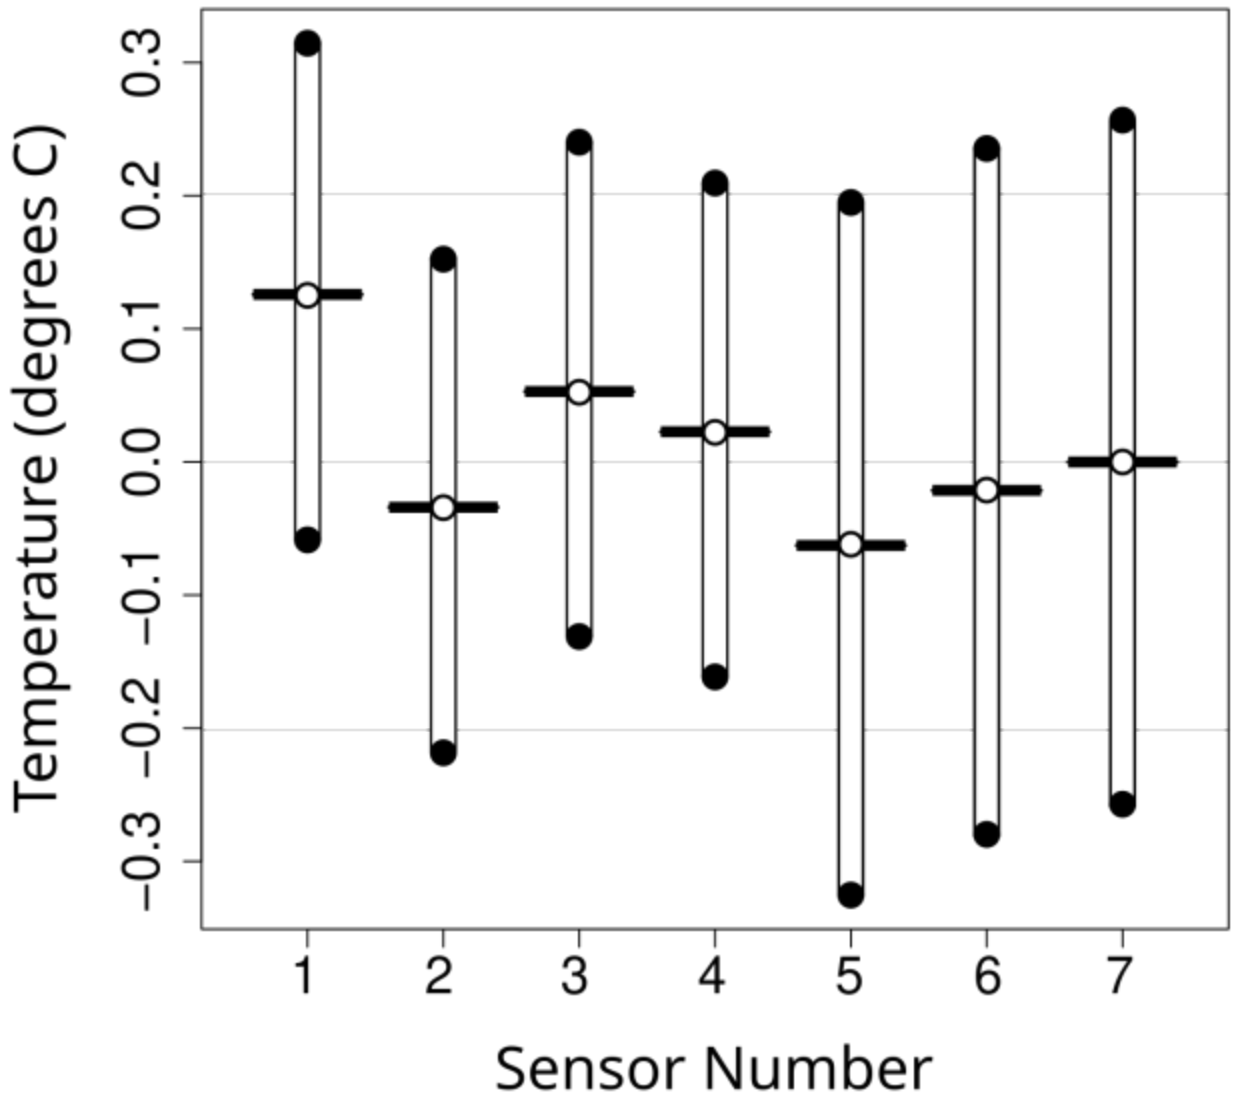

Figure 8. Posterior temperature means (black lines), medians (white circles), and 95% credible intervals (white bars and black circles) for each sensor node. Note that all sensor nodes have a CI of ~0.5C. The sensor nodes in the small, hyporheic side channel (1-4) have smaller credible intervals than those in fully connected areas of either the main channel or side channel (5-7).

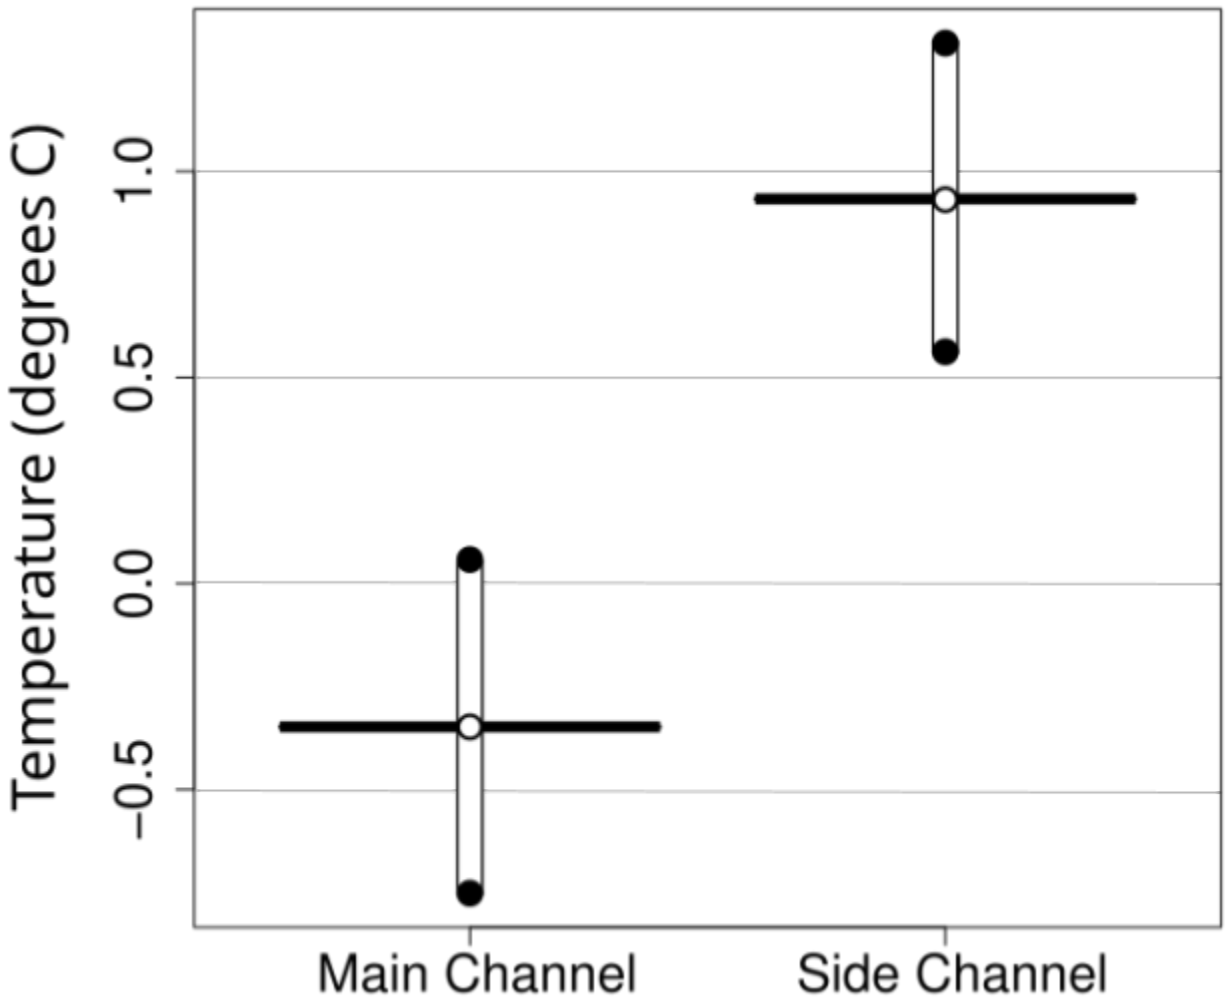

Figure 9. Model posterior means (black line), medians (white circles) and 95% credible intervals (white bars and black circles) of temperatures stratified by position in river, in the main channel or side channel.

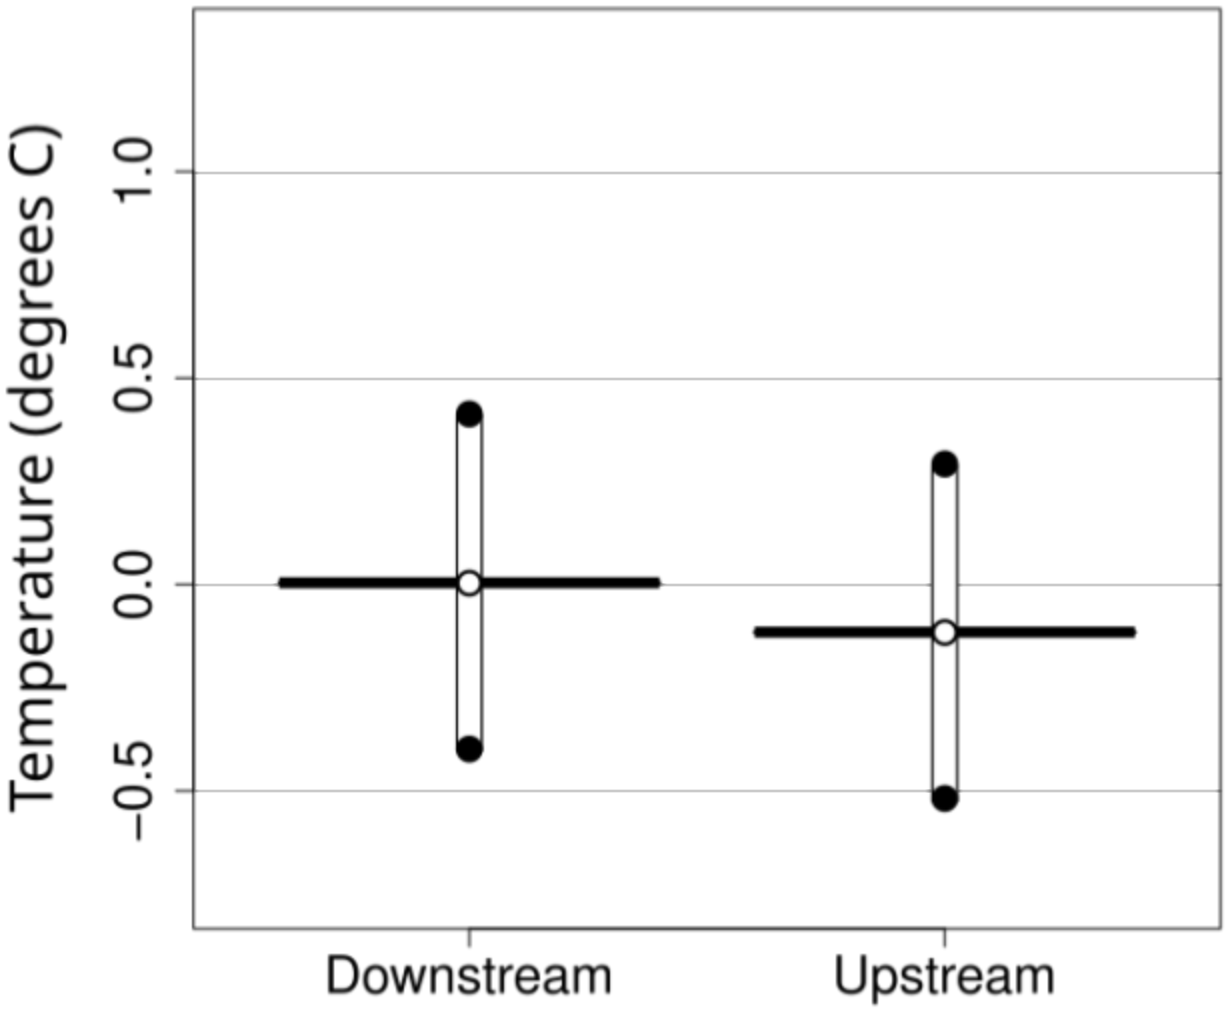

Figure 10. Model posterior means (black line), medians (white circles) and 95% credible intervals (white bars and black circles) of effect of position in the river above and below the merge-point in the river.

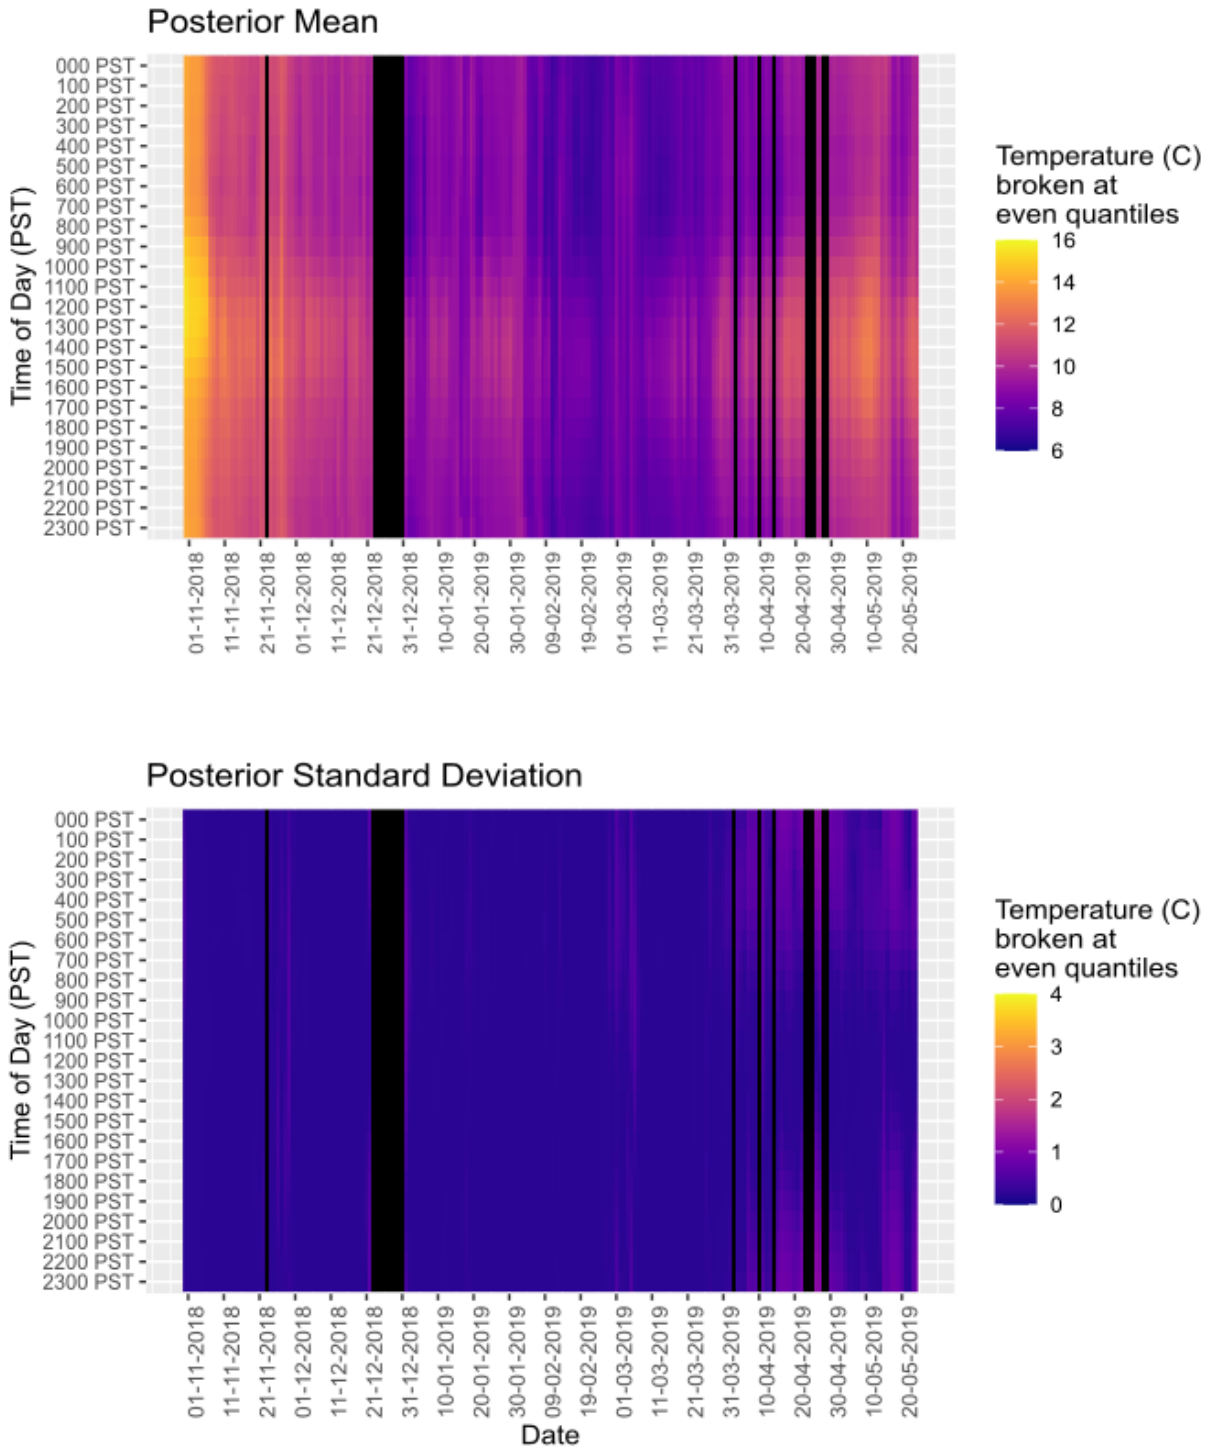

Figure 11. Model posterior mean (A) and standard deviation (B) temperatures (°C). Black bands correspond to absent data, or periods of extreme variability resulting from poor sampling during extended periods of severe inclement weather, which interfered with radio communications.

Seasonal and diel changes in spatial structures of sensed river temperatures

Abstract

Fish move through rivers, interacting with water temperatures that vary over space and time. At the same time, the waters flow downstream, yielding direct longitudinal connections within the channel. While seasonal and diel fluctuations are well established, there is little information on the spatial structures and coherences of these fluctuations. In this study, the primary research question addresses whether the longitudinal and lateral spatial structures of river temperatures are static or if they change over time, by time of day or day of year. Secondarily, we question if the time-dependence of these spatial structures is related to river flow rate and/or surface heating indexed by daily solar radiation. To identify the spatial structures, we generated semi-variograms both laterally (stream-orthogonal) and longitudinally (stream-oriented) through a perennially wetted side-channel of the lower Yuba River, a multi-threaded river in California's Central Valley. Temperature data were collected at approximately 15-minute intervals at 20 sensors from November 2018 to May 2019. The lateral and longitudinal distances between sensors were used to generate semi-variograms, calculated for specific time of day and specific day of year. We found that the spatial structures of fluvial temperatures were time varying and related this time dependence to river flow rate and daily solar radiation. Lateral structure in temperature (semi-variogram slope) reached a maximum in the middle of the night, a secondary maximum in the afternoon, and a minimum during the remainder of the day, which we relate to patterns in corresponding diel cycles in solar radiation. Lateral structure also varied over the seasons, which we relate to changes in river flow rate, being most pronounced in November-December 2018, and disappearing rapidly with the onset of river floods in January

2019. Longitudinal structure did not exhibit a coherent diel cycle, but we did observe seasonal changes in semi-variogram slope with marked longitudinal structure during low warm-water flows in November, high cold-water flows in February, and with the return of strong diurnal warming in May. These changes in flow modulate the level of connection between the side-channel and the mainstem of the river. River waters perpetually flow downstream, yielding questions about the permanency of patterns and structures within flowing waters.

Introduction

Fluvial temperature

In regulated rivers, the initial state of mean fluvial temperatures is set by the reservoir from which the river is fed. Intuitively, deeper reservoirs are colder, and water drawn from deeper within reservoirs is colder than water drawn from shallower in the reservoir. This initial state pertains more to mean, than to in-river temperature dynamics or spatial variations of fluvial temperatures. Spatial and temporal patterns in river temperatures are complex and influenced by several factors.

General spatial trends in fluvial temperature are well characterized: both mean and variance of temperatures are expected to increase from headwaters to mouth (Ward 1985). This trend is further discussed in Chapter 2, in which we address temporal variability. In this chapter, we focus on smaller-scale spatial patterns in temperature, i.e., patchiness. Insolation warms the river (Danehy et al 2005), but is mediated by shading from riparian vegetation (Johnson 2004). Seasonal inputs from precipitation and snow-melt cool the channel (Danehy et al 2005). Hyporheic flows cool small areas of the channel (Ebersole et al 2003), but are mediated by

geomorphic features (Bilby 1984). Finally, sediment can modulate water temperature through conduction (Story et al 2003).

The impact of solar insolation on river temperature is strong, but the effect is heterogeneous and localized (Jackson et al 2021). The warming effect of insolation is greatest when waters are slower moving and shallow (Daniels and Danner 2020). The warming associated with insolation was found to be mitigated by increased flow and snow-melt (Smith 1975, Beschta and Taylor 1988, Daniels and Danner 2020, Jackson et al 2021). River margins are particularly susceptible to warming from solar insolation as they are often shallower and closer to the substrate, which is warmed radiatively by sunlight. These areas are also characterized by slower-moving waters, yet their proximity to on-shore vegetation make them likelier to be shaded (Caissie 2006). As a result, they tend to demonstrate high thermal heterogeneity. Indeed, shading was found to have an extreme effect on local temperatures, particularly at channel margins (Beschta and Taylor 1988, Holtby 1988, Hetrick et al 1998, Johnson and Jones 2000, Johnson 2004, Caissie 2006). While shading effects meaningfully reduce local surface water temperatures, their impact on larger-scale stream temperatures is negligible, particularly when flows were higher (Hetrick et al 1998).

In Mediterranean climates, seasonal variations can be extreme (Andrew and Sauquet 2017, Cid et al 2017, Daniels and Danner 2020). In these systems, precipitative cooling (either by direct cooling like snow or cold precipitation, or by increased flow from headwaters) aligns temporally with seasonal minima in warming through radiation (solar), conduction (from shallow sediments warmed by solar radiation), and convection (contact with air), exacerbating them rather than mitigating them (as summer-wet climates might), yielding very cold winter temperatures, and warmer summer temperatures. Given the localized effects of precipitation, this may have implications for the patchiness of fluvial temperatures. Seasonal and interannual temperature

effects are difficult to predict in Mediterranean-climate rivers as the duration and magnitude of precipitation are more variable than in other humid climates (Andrew and Sauquet 2017, Cid et al 2017).

Hyporheic flow is the movement of river water through sediments, and this water tends to be colder than surface flows. Hyporheic flows are underground, and are therefore not warmed by exposure to sunlight or convective warming from air (Caissie 2006). Hyporheic connections can extend as far as kilometers from active channels and can include movement through floodplains at meander bends (Ebersole et al 2003, Uehlinger et al 2003). Hyporheic flows and their cooling effects are localized, and likely to enter the river channel in particular geomorphic landforms such as riffles, where they can draw colder groundwater seepage into the channel (White et al 1987). As a result, the pools behind riffles are frequently colder (Bilby 1984, White et al 1987). Multi-threaded rivers are more likely to have non-surface flows, which increases groundwater seepage (Bilby 1984, Ebersole et al 2003). Dominant hyporheic flows are associated with lower maximum temperatures and lower temperature variabilities (Ebersole et al 2003). Hyporheic flows are also mediated by sedimentary material: mud seals for example were found to be able to completely isolate groundwater from surface flows (Harrison and Clayton 1970). Aside from groundwater inputs and non-surface avoidance of insolation, hyporheic flows also have the greatest contact areas between sediment and water. This contact allows for conductive heat transfer (Johnson and Jones 2000, Story et al 2003). These conductive transfers increase as rivers become more threaded as the total surface area of contact between water and sediments increases. Differences in sediment also impacted rates of hyporheic flow and sediment conduction: waters flowing through alluvial sediments had less extreme maxima and smaller thermal fluctuations than did waters flowing over bedrock, which tended to increase thermal maxima, yielding wide ranges and fluctuations (Johnson and Jones 2000, Johnson 2004).

Study Area

The lower Yuba River (LYR) is a multi-threaded river in California's Central Valley. The river is constrained by engineered training-berms, has a gravel/cobble bed, and during the dry season is predominantly fed by low-level releases out of a deep, large reservoir, keeping the river colder than many other rivers in California's Central Valley in the dry season. The LYR merges with the Feather River near Marysville, CA. The Feather, in turn, flows downstream to the Sacramento River, which flows to the San Francisco Bay (Figure 1).

The study area was a ~3 km reach immediately downstream from Daguerre Point Dam, an ~8 m tall dam, built in the early 1900s to help block the downstream migration of hydraulic mining sediment and to divert water for nearby agriculture. One consequence of this dam is that the water is well-mixed in a plunge pool that then feeds the reach just below the dam (Figure 1). This reach was channelized and constrained into a mainstem southern channel and a higher-elevation, northern overflow channel by flanking and medial artificial training-berms. A series of floods in 2017 and 2018 scoured the overflow channel at the upstream end, re-engaging it with perennial flow. Floods also caused significant lateral migration of the mainstem to the degree that it entirely cut through the medial berm ~1 km downstream, thereby reconnecting the two branches here as well. The northern channel downstream of this merge point remains perennially hyporheically connected, but is only connected by surface flows during floods (Figure 1).

Study purpose

We installed 35 submerged temperature sensors from November of 2018 to May of 2019 (Burman and Gao et al 2023). A Bayesian model of observed temperatures revealed that the observed temperatures were best explained by an interaction between diel and day-of-sampling

(used as a proxy for seasonality) temporal effects (Chapter 2). This same model found little difference in mean temperature between upstream and downstream sensors, or between those in the mainstem and side channel. As temporal effects dominated the model of the full-reach, we anticipated that they might also dominate the spatial autocorrelation structures within our sensed temperature data.

In this chapter, we consider the 20 sensors in the contiguous, downstream side channel (described in the Field Methods section below) and pose the following research questions (RQ):

> **RQ1**: Is the spatial autocorrelation structure of temperature stable, or does it change over time (either diel or seasonal)?
>
> **RQ2**: Do temporal changes in spatial structure correspond to changes in river flow and/or radiative surface heating?

We investigated these questions using spatial semi-variograms and hypothesized that the semi-variograms would be time-dependent. In particular, given the strong, localized effects of insolation and shading (Beschta and Taylor 1988, Holtby 1988, Hetrick et al 1998, Johnson and Jones 2000, Johnson 2004, Danehy et al 2005, Caissie 2006), we hypothesized that the longitudinal and lateral semi-variograms would have the greatest slopes shortly after the solar maximum around noon, and remain high throughout the period when water temperatures were warmest (around 2pm: Burman and Gao et al 2023). We anticipated that these slopes would then decrease until the river was relatively homogeneous around midnight, when solar radiation was lowest, and remain low through the period when river temperatures were coldest (around 2 am: Burman and Gao et al 2023). We also calculated semi-variograms over 20-day intervals throughout the sampling period and anticipated that semi-variograms would change seasonally such that when flows rose and connected surface flows in the side channel and mainstem, this

would homogenize temperatures. We anticipated semi-variogram slopes, both lateral and longitudinal would decrease rapidly when river flows increase in winter, and remain low through the high-flow season

Scientific contributions

By generating time-varying semi-variograms of water temperatures both laterally and longitudinally in a perennially wetted side channel, we show that spatial structures were not time-invariant. Spatial structures changed following both diel and seasonal signatures. We compare these changes in spatial structure to the median hourly solar radiation, the daily total solar radiation, the 95% inter-quantile range of flow, and daily mean flow. From these comparisons, we infer that spatial structures in the side channel were primarily controlled by both flow and solar radiation. Notably, diel patterns in lateral spatial structure were related to changes in radiative warming, and both lateral and longitudinal patterns in seasonal structure were related to increasing flows that reconnected the surface flows in the side channel to the mainstem of the river.

Methods

Field methods

To study fluvial temperatures within this reach, we deployed a wireless sensor network consisting of 35 temperature sensors located in both the primary and side-channels both above and below a merge-point in the middle of the study reach (see Study Area section; Figure 1). As river temperatures in wide gravel and cobble rivers can vary more across a channel than longitudinally (Caissie 2006), associated with diverse morphological units and lateral groundwater discharges from hillsides, it was imperative to measure temperatures at different

locations across the wetted channel (laterally) as well as along the channel (longitudinally) and at a suitable frequency in time (temporally). In this analysis, we use only the 20 sensors in the downstream side channel (Figure 1), which was a contiguous body of water, and for discharges below ~141 m^3/s, was exclusively hyporheically fed. The wireless sensor network and data collection and processing are described in Burman et al (2020) and Burman and Gao et al (2023). There were five sensors on each cross-channel transect, which were attached to the sensor nodes that contained all of the necessary electronics and wireless communication hardware. Water temperature was measured about four times per hour (sensor nodes did not have real-time clocks so sampling frequency was temperature dependent, fluctuating between about 14 and 16 minutes).

The field methods used to address question one involved measuring temperatures multiple times per hour. Each sensor transect sampled the associated five sensors about four times per hour. Further, given strong seasonal changes in river temperatures, surface warming, and river flow rates, we planned to sample for a year to capture these seasonal fluctuations. Unfortunately, our field study was terminated prematurely by a river enhancement project (Southall et al 2022).

Here we focus on the four cross-channel (lateral) transects that were placed in the side channel (Figure 1) and provided data for the duration of the study (November 2018 to May 2019):

- Transect 1 (sensors 11-15) traversed a deep pool with clear water at the downstream end of the side channel. Shallow parts of the pool contained aquatic vegetation. The pool was perpetually wetted, and had a bar at the downstream end that acted like a check valve, preventing flow from the mainstem into the pool. The pool was ~ 50 m across. The sensors installed were spaced ~ 10 m apart.

- Transect 2 (sensors 21-25) was in a riffle just upstream of the deep pool (spanned by transect 1). These sensors were spaced 5 m apart. The furthest upstream sensor on the string was placed in an extremely cold, but very small stream that was generally disconnected from the rest of the side channel and entirely hyporheically fed. This tiny channel then fed into a riffle in the larger side channel. The second to last sensor was in the deepest part of the pool in the riffle-pool couplet, and the final sensor was on a shallow bar.

- Transect 3 (sensors 31-35) was in a glide within the side channel. Near the northern channel margin there was a small pool and at the other end (in the deepest part) was a slightly faster glide. Sensors were spaced 5 m apart.

- Transect 4 (sensors 41-45) was in a very small riffle that was very slow moving during summer baseflow and was often just a few centimeters in depth. Just upstream from this small riffle was a large pool that fed the side-channel. Sensors were spaced about 5 m apart.

Some sensors were damaged during the study, remaining online until February 2019, while others provided data through May 2019.

Raw temperature, hydrologic, and solar radiation data

Temperature data from all sites were reported in Burman and Gao et al (2023). Here we consider a subset of the installed sensors - the 20 sensors in the side channel (Figure 1). For each sensor we calculate the daily mean temperature to track spatial and temporal variability and the seasonal cycle, and we calculate the daily 95% inter-quantile range (i.e: 97.5th quantile - 2.5th quantile) as a measure of the strength of the diel cycle.

River flow rate data were acquired from a USGS flow gauge about 3.5 km west (downstream) of the downstream end of the study reach (USGS 2022). Though most of the measured water flows through the mainstem, with only a small portion flowing through the side channel considered in this study, we assume that flow in the side channel is correlated with that in the main stem. Units were converted from ft^3/s to m^3/s and we calculate daily mean flow and a daily 95% inter-quantile range (97.5th quantile - 2.5th quantile).

Finally, we acquired solar radiation data from the National Solar Radiation Data Base (Sengupta et al 2018, NSRDB 2022) as an index of surface warming, i.e., to differentiate between long and sunny days from shorter or cloudier days. We used the global horizontal irradiance (GHI), which is a calculated metric that incorporates both the direct solar radiation (Direct Normal Irradiance: DNI) and the solar radiation reflected or refracted by the atmosphere (Diffuse Horizontal Irradiance: DHI) to create a metric of the total power of sunlight hitting a square meter of the earth's surface ($\frac{W}{m^2}$, Sengupta et al 2018). DNI and DHI (and by extension, GHI) are calculated by the NSRDB using remotely sensed data from NOAA and NASA satellites; these models are validated and calibrated using a network of strategically placed pyrheliometers and pyranometers (Sengupta et al 2018). The data were queried for the time-period during which our sensors were deployed (01 Nov 2018 - 01 Jun 2019), at Beale Air Force Base, which is about 6 km southeast of our study sites (NSRDB 2022). We summed the hourly GHI data for each day, yielding a new time series of the total radiation energy hitting a square meter near the river for each day in the study period ($\frac{W}{m^2} * hr$, Sengupta et al 2018, NSRDB 2022).

Analytical methods

Given that LYR temperatures varied over the course of the day, and across sampled days (Burman and Gao et al 2023, Chapter 2), we sought to identify the underlying spatial scales on

which differences in these temporal effects were observed. We created semi-variogram clouds both laterally and longitudinally, calculated for different sampling days and for different times of day. We compared the slopes in these semi-variograms (an index of spatial autocorrelation structure) with cumulative daily solar radiation and the hourly median solar radiation (GHI; NSRDB 2022), as well as both daily mean and daily 95% inter-quantile range of flows as measured just below the study reach (USGS 2022) over the same time period.

Sensed temperature data were first separated into 12 bins of two hours each, starting with midnight to two am (PST), then two am to four am and so on through the full 24-hour day. Further, sampling days were divided into 11 bins of 20 days each. Once these bins were defined, we created lateral and longitudinal semi-variogram datasets. We first created lateral semi-variogram datasets by randomly selecting 20,000 sensor readings (with replacement) for each of the twelve 2-hour blocks, and for each of the eleven 20-day blocks. For each of these 20,000 sensor readings, we selected a paired reading from the same sensor string (same transect) within the same two-hour block and the same 20-day block. We therefore created, for each of these temporal blocks a dataset of 20,000 pairs of sensor locations and sensor readings. These paired sensor readings were then each used to compute a semi-variance, and plotted against the distance between the paired readings. This was repeated for each 2-hour or 20-day time block. Given that channel widths varied, we computed the sensor spacing as a decimal of the transect length (20 m for 5-m sensor spacing or 40 m for 10-m sensor spacing).

This process was repeated to generate the longitudinal semi-variogram datasets. We measured the stream-oriented downstream distance of each sensor string from a fixed point near the top of the side-channel. We again took 20,000 samples per 2-hour block and 20-day block. This time, we placed no geographic restrictions on the paired sensor readings, they could come from

any of the sensors in the river. We then used the difference in stream-oriented downstream distances (i.e., longitudinal distance) to compute geographic distances between sensor pairs.

(Semi-) Variance, Variogram, and Cloud

Semi-variance, is a metric of dissimilarity between values at two or more positions in space (Cressie 1993, Wackernagel 1998): larger semi-variance indicates more dissimilarity. By plotting this dissimilarity against the geographic separation between two points, we can quantify how the dissimilarity changes over spatial distances (O'Sullivan and Unwin 2010). When generated from pairwise comparisons of values derived from experimental data, these are known as variogram clouds (Wackernagel 1998). Notably, the terms variogram and semi-variogram describe the same concepts and, we treat them as interchangeable as is common in the literature - although not all authors do so (see Bachmaler and Backes 2008 for a more exhaustive consideration of this terminology). For two points (i and j), with sensor readings t_i and t_j separated by distance $d_{i,j}$, the experimental semi-variance ($\bar{\gamma}$) is computed by:

$$\bar{\gamma}(d_{i,j}) = \frac{(t_i - t_j)^2}{2} \quad (1)$$

with cross-channel decimal distance (d) for sensor spacing (s_S) defined as:

$$d_{i,j} = \frac{\sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}}{4 \times S_S} \quad (2)$$

Longitudinal downstream distances (d) were computed more simply, where the downstream distances of the two points (p_i,p_j) is the difference between their respective stream-oriented downstream distances from a fixed point (p_0). Because all distances are relative, the precise location of this point is arbitrary.

$$d_{i,j} = abs((p_i - p_0) - (p_j - p_0)) \quad (3)$$

In order to create a single, cohesive curve from these semi-variance-distance point clouds, past work has often used either mean variance values, using either absolute distance values, or mean values within binned geographic distances called lags (Wackernagel 1998). We opted to instead use a linear model smoothing kernel to create a smooth line for each of the (twelve 2-hour and eleven 20-day) time-period bins considered, for both orientations (lateral and longitudinal). Semi-variances that increase with spatial separation are said to be positively autocorrelated: proximate point-pairs are more similar than are distance point-pairs. Semi-variances that decrease with spatial separation are said to be negatively autocorrelated: proximate point-pairs are less similar than distant point-pairs. Semi-variances that neither decrease nor increase with spatial separation are said to not be autocorrelated (O'Sullivan and Unwin 2010), i.e., there is no spatial structure. Idealized variograms that reflect autocorrelation follow an asymptotic curve, with semi-variance increasing with distance from the considered point, then reaching a threshold beyond which, the points are not affected by local autocorrelation, at which point the semi-variance flattens to an asymptote (O'Sullivan and Unwin 2010).

Question testing

For both the diel cycle (2-hour variograms) and the seasonal cycle (20-day variograms), we compared the slopes of the plotted variograms for both longitudinal and lateral pairs and plotted these slopes against time of day (for 2-hour variograms) or day of year (for 20-day variograms). The slope is a measure of spatial dissimilarity and the time-dependence of the slope reflects development and dissipation of spatial structure. If the slopes of the variograms are approximately the same over time (i.e., slope values fall within the standard error bands), then we can conclude that autocorrelation structures are time invariant. To contextualize these

time-of-day or day-of-year changes, we compared these results to time series data on solar radiation (GHI; NSRDB 2022) river flow rate (USGS 2022) measured at nearby locations.

Results

In plots of the temperature data from all 20 sensors installed in the side channel of the study reach, water temperature in the side channel decreased through the fall and early winter, from a maximum of about 16^{o}C in early November 2018 to a minimum of about 8^{o}C in late February 2019 before increasing again in late winter and spring to about 12^{o}C in May 2019 (Figure 2b). In addition to the seasonal signal, short-term fluctuations are evident, with temperature changing as much as 2^{o}C over a few days, and differences between sensors (excepting sensor 21 – see below) can be as much as 2^{o}C, specifically during the cooling phase in November and December. The diel range in water temperature varies from 0 to 4^{o}C, without a clear seasonal signal (Figure 2a). The diel signal is persistently low for a couple of weeks in late February and early March, and then strong for a couple of weeks in mid-March, but in general the diel signal can be as strong in winter as in fall and spring.

Certain sensors experience larger temperature swings than do others (Figure 2a,b; Burman and Gao et al 2023). Notably, sensor string two (sensors 21-25), which traverses a massive riffle, a very cold hyporheically fed backwater, and a bar, is subject to the greatest temperature variability (Figure 2a,b). Sensor 21, within sensor node two is placed on a shallow bar that, while wetted at all times of observation, was often in a small amount of standing water, subject to little or no flow. Many other sensor strings were quite consistent laterally. Sensor node one (sensors 11-15) for example, which spans the width of a large lagoon, is quite homogeneous, as is sensor node four (sensors 41-45), which traverses a relatively well-mixed riffle at the upstream extent of the study reach (Figures 1,2; Burman and Gao et al 2023, Chapter 2).

As we expect daily mean temperature and diel variability to be modulated by the flow rate in the river and by the strength of surface heating, we plot these data for the same period as the water temperature data (Figure 3). Daily mean flow rate was consistently low in November and December 2018 (< 50 m^3/s) and rose abruptly in January with flow pulses exceeding 100 m^3/s; Figure 3b). Maximum flows of 600 m^3/s occurred during pulses in February and March. During this time period, the daily 95% inter-quantile range reached ~ 500 m^3/s during storm events in February and March (Figure 3a) showing that these peaks are consistent with short-duration flood events. In late March, mean flows decreased and remained consistent until June 2019 (Figure 3b). Cumulative daily radiation shows a clear seasonal cycle (Figure 3c), dropping to a minimum clear-sky value of about $120\ \frac{W}{m^2} * hour$ in December (shortest days) and reaching a maximum of about $360\ \frac{W}{m^2} * hour$ in June (longest days). However, on many days the sky is not clear and GHI is notably less, on a day-by-day basis (Figure 3c). Notably, local minima in solar radiation (consistent with periods of high precipitation) tend to correspond with periods of high flow; these minima often precede the local maxima in flow by a day or two (Figure 3b,c; NSRDB 2022, USGS 2022), e.g., late February 2019. When grouped by hour of day (PST), we find that median solar radiation (GHI; NSRDB 2022) is zero ($\frac{W}{m^2}$) overnight, and begins to rise linearly around 5 or 6 am, reaching a maximum around 11am at about 750 ($\frac{W}{m^2}$) , which remains until the noon hour, at which point median solar radiation falls linearly until again returning to zero around 6pm, and remaining zero until 5 or 6 am the following morning (Figure 4, NSRDB 2022). Even in winter, water temperature increased during sunny/warm weather with low river flow, e.g., late January - and marked spatial differences emerge in both mean temperature and diel range (Figure 2).

We use semi-variograms to quantify spatial structure (both lateral and longitudinal) and to quantify both diel and seasonal fluctuations in this spatial structure. Using semi-variograms to assess lateral spatial structure, across all 12 two-hour blocks, we see an increase in semi-variance as the distance between compared sensor points increases (Figure 5), which is consistent with positive autocorrelation. These rates are not homogeneous through time (Figure 5). The slope of these effects reaches a maximum near the minimum daily temperature, around 2 am, and a minimum after the 2 pm daily maximum temperature, around 5 or 6 pm. Further, there is a slight increase in the afternoon between around noon and 4 pm (Figure 5), which corresponds with the daily maximum temperature. Considering the lateral semi-variances, based on a 20-day sampling period, slopes are near zero for most winter sampling periods (Figure 6). There exist positive slopes only for the sampling windows during November and December (Figure 6). These correspond with the late autumn and early winter months, when local temperatures are falling.

When looking at longitudinal temperature profiles, the result was not nearly as consistent. Over the course of a day, among the slopes of the semi-variance curves, we observed multiple local maxima and minima - which are all weak (Figure 7). Notably, in addition to the trend, with a global maximum near the daily minimum temperature (2 am) and a mid-afternoon peak, we also observed local maxima in the early morning around 6-8 am (Figure 7). Looking at the longitudinal sampling relative to the 20-day sampling period, we see a trend of high auto-correlation in the first 20-day period, falling quickly in late November, then rising to a local maxima in February, before falling again, reaching a local minimum in April, and finally, rising again in May (Figure 8).

Discussion

Our analyses of temperature data sampled at high-frequency from four sites along multiple transects in a side channel of the LYR build upon the previously modeled diel and seasonal cycles in observed temperatures from the same study reach (Chapter 2). The analyses in the present study found a diel signature in lateral spatial structure of temperatures. Further, there exists a seasonal trend in both lateral and longitudinal temperature spatial structures. Broadly speaking, diel patterns aligned well with trends in solar radiation, and seasonal patterns aligned well with patterns in flow. Given the observed interactive effects of seasonal and diel signatures (Chapter 2), we offer this as an avenue for future study: identifying the interactive effects on spatial structure of the two dominant time scales identified in chapter 2.

Spatial patterns vary over time (RQ1)

We observed differences between spatial autocorrelation structures both laterally and longitudinally, and for both the 2-hour and 20-day time bins. In the 2-hour time bins – of diel spatial structures – we anticipated that fluvial temperatures would be “patchiest” – with the most extreme positive autocorrelation structure (observed as a variogram with the highest slope) – in the afternoon, because insolation would warm shallow, slower-moving areas, mediated by local vegetative shading. This was not what we observed: while we found a local maximum in spatial autocorrelation structure in the afternoon, the global maximum was actually observed around midnight to 2 am, when water temperatures were coldest (Figures 2,5).

When considering lateral temperature variograms as delineated using the 20-day bins, we observed a single overriding trend: there is very little spatial autocorrelation for much for the sampling period, with much stronger autocorrelation structures in November and December, which then fell considerably in January, and remained low through the end of the sampling

period in May (Figure 6). The onset of this break-down in cross-channel autocorrelation corresponded with the coldest months in the study reach (Figure 2). Because of the Mediterranean climate of California's Central Valley, this is the wettest period as well. Precipitation could affect river temperatures locally, thereby increasing spatial heterogeneity – especially in Mediterranean climates.

Longitudinal profiles were a bit more difficult to disentangle. The slopes in question followed no consistent trend when considered through the lens of time of day (using the 2-hour bins), we see that nearly all of the slopes are close to zero, and are similar to each other (Figure 7). Perhaps the best interpretation of this figure is that the diel effects of temperature are best observed across the channel and in proximate regions of a reach. And, when considered over longer, longitudinal spatial scales, the persistent downstream movement of rivers mitigates the localized effects of diel signatures.

Longitudinal changes in river temperature were much more pronounced when considering the temperature signatures delineated by the 20-day bins. We observe the patchiest autocorrelation structures in November (Figure 8). The decay of the longitudinal autocorrelation signature occurred in late November to early December, nearly a month earlier than observed in the lateral autocorrelation signature. For much of the remainder of the sampling period, we observed minimal longitudinal autocorrelation, aside from a local maximum in early February (Figure 8). This was unexpected as we anticipated that winter precipitation could increase longitudinal spatial heterogeneity.

Relating these patterns in spatial structure to surface heating (GHI) and river flow (RQ2)

The observed pattern of diel lateral structure reached a local maximum around 2 pm, with a global maximum around 2 am, and global minima around noon and 6-8 pm (Figure 5). The afternoon local maximum coincided with the maximum solar radiation (Figure 4), around noon and slowly decreasing over the afternoon (Figure 5). What we did not anticipate was that lateral autocorrelation would be greatest when river temperatures were lowest, reaching a maximum at the daily temperature minimum in the middle of the night, when solar radiation was zero (Figure 4; Figure 5). We posit that a possible explanation for this pertains to the complexity of temperature interactions, there are localized cold-regions generated by things like hyporheic flows, geomorphology, groundwater inputs, and precipitation (Beschta and Taylor 1988, Holtby 1988, Hetrick et al 1998, Johnson and Jones 2000, Johnson 2004, Caissie 2006, Daniels and Danner 2020, Jackson et al 2021). Perhaps in a smaller side channel such as the study region, solar insolation simply overrides these other heterogeneous inputs, and homogenizes the river across the channel. In such a model, then at night, without these homogenizing solar warming effects, the other, local effects would dominate, yielding a maximum in semi-variance when solar radiation was at a minimum.

Considered seasonally (20-day time bins), lateral spatial structure had a maximum semi-variogram slope in November and December 2018, falling rapidly in late December and into January 2019 (Figure 5). In late 2018, we observed that flows were very low, and relatively homogeneous (Figure 3): well below the bankfull flow of the LYR ($\sim 141\ \frac{m^3}{s}$). Both mean flow and flow 95% inter-quantile range increased in early 2019 (Figure 3). The low flows observed in late 2018 would have meant that there were no surface-flow connection between the side

channel and the mainstem of the river, and instead the side channel was entirely hyporheically fed (Figures 1,3). In early 2019 when flows increased, particularly when they exceeded the bankfull flow ($\sim 141\ \frac{m^3}{s}$), the side channel would have been re-connected with the mainstem by surface flows (Figures 1,3). We posit that this was likely the cause of the observed breakdown in lateral spatial structure: later in the winter, the flows were sufficient to actively reconnect the mainstem with the side channel, and these flows in turn homogenized and mixed the side channel, eliminating the strong lateral spatial structures observed in late 2018 (Figure 6).

The seasonal (20-day time bins) changes in longitudinal temperature spatial structure exhibited a global maximum in late 2018, falling precipitously in December 2018 and into January 2019 (such that there was in fact no autocorrelation in 2019) aside from a small increase in heterogeneity in early February (Figure 8). At the end of 2018, both mean daily flow and 95% inter-quantile range in flow were relatively consistent and low (Figure 3a,b), and likely resulted in the high spatial heterogeneity in longitudinal temperatures (Figure 8). This was similar to the observed trend in the seasonal, lateral semi-variograms (Figure 6). The lateral and longitudinal seasonal semi-variograms differ in that the longitudinal semi-variogram slopes fall much more rapidly in late November, then rise starting in January to reach a secondary maximum in February before falling again through the end of the sampling period (Figure 8). The fall in spatial autocorrelation in late November or early December 2018 corresponds with intermittent days of diminished daily total GHI (Figure 3c). These lower GHI days (likely caused by overcast skies and precipitation) seem to have coincided with a decrease in longitudinal warming of the side channel, which in turn caused high semi-variance slope observed in early November and subsequent fall in early December (Figure 3,8). Further, the secondary maximum in semi-variogram slope seen in late February corresponds to a wetter period, with greater connections to the mainstem, and greater variability in daily solar radiation (Figure 3a,b,c). Notably, the connections to the mainstem had different effects longitudinally than laterally.

While this connection thoroughly mixed the side channel laterally, this seemed not to mitigate the longitudinal spatial structures. At the downstream end of the study reach, at sensor node one, there was a large pond (Figure 1). This pond was nearly 500 m in length, about 250 m wide, and workers could not stand in the middle (at least 2 m deep). There was a bar below this pond, which provided stage control and acted something like a check valve, only allowing flow from side channel through the pond into the mainstem, but limiting flow from the mainstem into the pond (as observed by field workers). As such, there was a long residency period for waters in this pond, with greater opportunities for surface warming, and less opportunities for mixing via higher flows. This pond likely acted as a fairly consistent thermal mass, regardless of precipitation, insolation, or flows, for any longitudinal considerations of spatial structure. We found that solar radiation and river flows likely interacted with the temperatures of a perennially wetted, hyporheically connected river side channel. Further, over the longer distances of the longitudinal semi-variograms, it was possible that direct precipitation inputs yielded thermal heterogeneity, which were not stratified laterally, where distances were smaller.

Future study and contributions

There are many opportunities for further study of semi-variance of fluvial structures, particularly temperatures. As one might expect, fluvial temperatures are not stationary as they are not isotropic (we found very different semi-variogram slopes in our lateral and longitudinal models). The consistency of these models, and the strong indications of patterns in their change over time suggest that variances in fluvial temperatures of flowing waters are a manifestation of external forcings, instead of being predicated by an originating state. They are warmed and vary in accordance with environmental factors around the river. Rivers do not simply transport flowing waters, the waters are changed during the process of flowing. While the water may be moving and ephemeral, these forcings in fact have structures that can persist, and yield meaningful

patterns that can be interpretable. We posit that there exist potentials to identify underlying forcings and inputs using the underlying structures that manifest in observable fluvial states, specifically temperatures. The spatial structures of motile systems like rivers yield exciting opportunities for further study. Few systems exist within a substrate in consistent, directed (and directional) motion.

Given the strong observed temporal effects on spatial structure, we also highlight that such systems require models of spatial structure that interact with temporal effects. Taking for example, the model in Chapter 2, we highlight the potential to extend this model with a non-stationary, anisometric spatial effect, which interacts with the temporal effects described in that paper. Dimensional, autoregressive terms seem to be effective in characterizing fluvial temperatures, potentially dominating mechanistic model effects like GHI, flow, or shading.

Acknowledgements

This work was supported by a cooperative agreement (F16AC00735) with the Anadromous Fisheries Restoration Program within the US Fish and Wildlife Service with the support of Dr. Beth Campbell and Paul Cadrett of the USFW. The lower Yuba River field crew from the Pacific States Marine Fisheries Commission were crucial in completing the relevant fieldwork, notably Loren Stearman, Duane Massa, and John Cleveland. Chelsea Martinez's contribution to this work cannot be overstated, this project would have failed without her diligent work, impeccable organization, and reliable field support in a challenging working environment. Richard Hambrick was instrumental in getting the development of network software off the ground. This work was taken over by Jingya Gao, a brilliant undergraduate student in computer science who worked independently and collaboratively to bring the networking software to fruition. Megan Werdmuller von Elgg provided crucial editorial support in the drafting of this document. Finally,

the folks who helped with this mountain of field work, without whom, the project could never have succeeded: Darcy Bostic, Annelise Del Rio, Arielle Gervais, Julius Henry, Vanessa Lo, Sean Luis, Seanna McLaughlin, Brianna Ordung, Lea Pollack, Kathy Russel, Rebecca Walker, Jason Wiener, Katie Woodworth, and Ken Zillig.

Literature Cited

Andrew JT, Sauquet E. 2017. Climate Change Impacts and Water Management Adaptation in Two Mediterranean-Climate Watersheds: Learning from the Durance and Sacramento Rivers. Water.9.126: doi:10.3390/w9020126

Bachmaier M, Backes M. 2008. Variogram or semivariogram? Understanding the variances in a variogram. Precision Agric. 9:173-175

Beschta RL, Taylor RL. 1988. Stream temperature increases and land use in a forested Oregon watershed. J Amer Water Resour As. 24:19–25. doi:10.1111/j.1752-1688.1988.tb00875.x.

Bilby RE. 1984. Characteristics and frequency of cool-water areas in a western Washington stream. J Freshwater Ecol. 2:593–602. doi:10.1080/02705060.1984.9664642.

Burman SG, Gao J, Ghosal D, Pasternack G, Fangue N. 2020. Design, Implementation, and Deployment of TempMesh: A Wireless Mesh Network to Aggregate River-Temperature Data. 2020 IEEE International Conference on Sensing, Communication and Networking (SECON Workshops). 10.1109/SECONWorkshops50264.2020.9149773.

Burman SG, Gao J, Pasternack GB, Fangue NA, Cadrett P, Campbell E, Ghosal D. 2023. TempMesh – A Flexible Wireless Sensor Network for Monitoring River Temperatures. ACM Trans Sens Netw. https://doi.org/10.1145/3542697

Caissie D. 2006. The thermal regime of rivers: a review. Freshwater Biol. 51:1389–1406. doi:10.1111/j.1365-2427.2006.01597.x.

Cid N, Bonada N, Carlson SM, Grantham TE, Gasith A, Resh VH. 2017. High Variability Is a Defining Component of Mediterranean-Climate Rivers and Their Biota. Water. 9(52).doi.org/10.3390/w9010052

Cressie NAC. 1993. Statistics for Spatial Data. John Wiley & Sons, Inc. DOI:10.1002/9781119115151

Danehy RJ, Colson CG, Parrett KB, Duke SD. 2005. Patterns and sources of thermal heterogeneity in small mountain streams within a forested setting. Forest Ecol Manage. 208:287–302. doi:10.1016/j.foreco.2004.12.006.

Daniels ME, Danner EM. 2020. The drivers of river temperatures below a large dam. Water Resour Res. 56. doi.org/10.1029/2019WR026751

Ebersole JL, Liss WJ, Frissell CA. 2003. Cold water patches in warm streams: physicochemical characteristics and the influence of shading. J Amer Water Resour As. 39:355–368. doi:10.1111/j.1752-1688.2003.tb04390.x.

Harrison SS, Clayton L. 1970. Effects of ground-water seepage on fluvial process. Geol Soc Amer Bull. 81:1217. doi:10.1130/0016-7606(1970)81[1217:EOGSOF]2.0.CO;2.

Hetrick NJ, Brusven MA, Meehan WR, Bjornn TC. 1998. Changes in solar input, water temperature, periphyton accumulation, and allochthonous input and storage after canopy removal along two small salmon streams in southeast Alaska. Trans Amer Fish Soc. 127:859–875. doi:10.1577/1548-8659(1998)127<0859:CISIWT>2.0.CO;2.

Holtby LB. 1988. Effects of logging on stream temperatures in Carnation Creek British Columbia, and associated impacts on the Coho salmon (oncorhynchus kisutch). Can J Fish Aquat Sci. 45:502–515. doi:10.1139/f88-060.

Jackson FL, Hannah DM, Ouellet V, Malcolm IA. 2021. A deterministic river temperature model to prioritize management of riparian woodlands to reduce summer maximum river temperatures. Hydrol Processes. doi.org/10.1002/hyp.14314

Johnson SL. 2004. Factors influencing stream temperatures in small streams: substrate effects and a shading experiment. Can J Fish Aquat Sci. 61:913–923. doi:10.1139/f04-040.

Johnson SL, Jones JA. 2000. Stream temperature responses to forest harvest and debris flows in western Cascades, Oregon. Can J Fish Aquat Sci. 57:30–39. doi:10.1139/cjfas-57-S2-30.

NSRDB (National Solar Radiation Data Base). 2022. National Renewable Energy Laboratory (NREL): Accessed 25 Dec 2022. Acquired data for 2018 and 2019: GHI, DNI, DHI at Beale Air Station: (Location code: 138262; lat: 39.13, long: -121.42)

O'Sullivan D, Unwin DJ. 2010. Geographic Information Analysis. Wiley: Hoboken, USA.

Sengupta M, Xie Y, Lopez A, Habte A, Maclaurin G, Shelby J. 2018. The National Solar Radiation Data Base (NSRDB). Renew Sust Energ Rev. https://doi.org/10.1016/j.rser.2018.03.003

Smith K. 1975. Water temperature variations within a major river stream. Nord Hydrol. 6:155-169

Southall N, Sawyer A, Diaz S, Hammersmark C, Bowles C. 2022. Golden opportunity. Civil Eng. 92 (2): 50-59. https://ascelibrary.org/doi/epdf/10.1061/ciegag.0001610

Story A, Moore RD, Macdonald JS. 2003. Stream temperatures in two shaded reaches below cutblocks and logging roads: downstream cooling linked to subsurface hydrology. Can J Forest Res. 33:1383–1396. doi:10.1139/x03-087

Uehlinger U, Malard F, Ward JV. 2003. Thermal patterns in the surface waters of a glacial river corridor (Val Roseg, Switzerland). Freshwater Biol. 48:284–300. doi:10.1046/j.1365-2427.2003.01000.x.

USGS. 2022. Yuba R NR Marysville. [https://waterdata.usgs.gov/monitoring-location/11421000/#parameterCode=00060&startDT=2018-11-01&endDT=2019-06-01]. [accessed: 10 Dec 2022]

Wackernagel H. 1998. 6: Variogram cloud. in: Multivariate Geostatistics. Springer-Verlag: Verlin, Germany.

Ward JV. 1985. Thermal characteristics of running waters. In: Davies BR, Walmsley RD, editors. Perspectives in Southern Hemisphere Limnology. Dordrecht: Springer Netherlands. p.31–46.

White DS, Elzinga CH, Hendricks SP. 1987. Temperature patterns within the hyporheic zone of a northern Michigan River. J N Amer Benthol Soc. 6:85–91. doi:10.2307/1467218

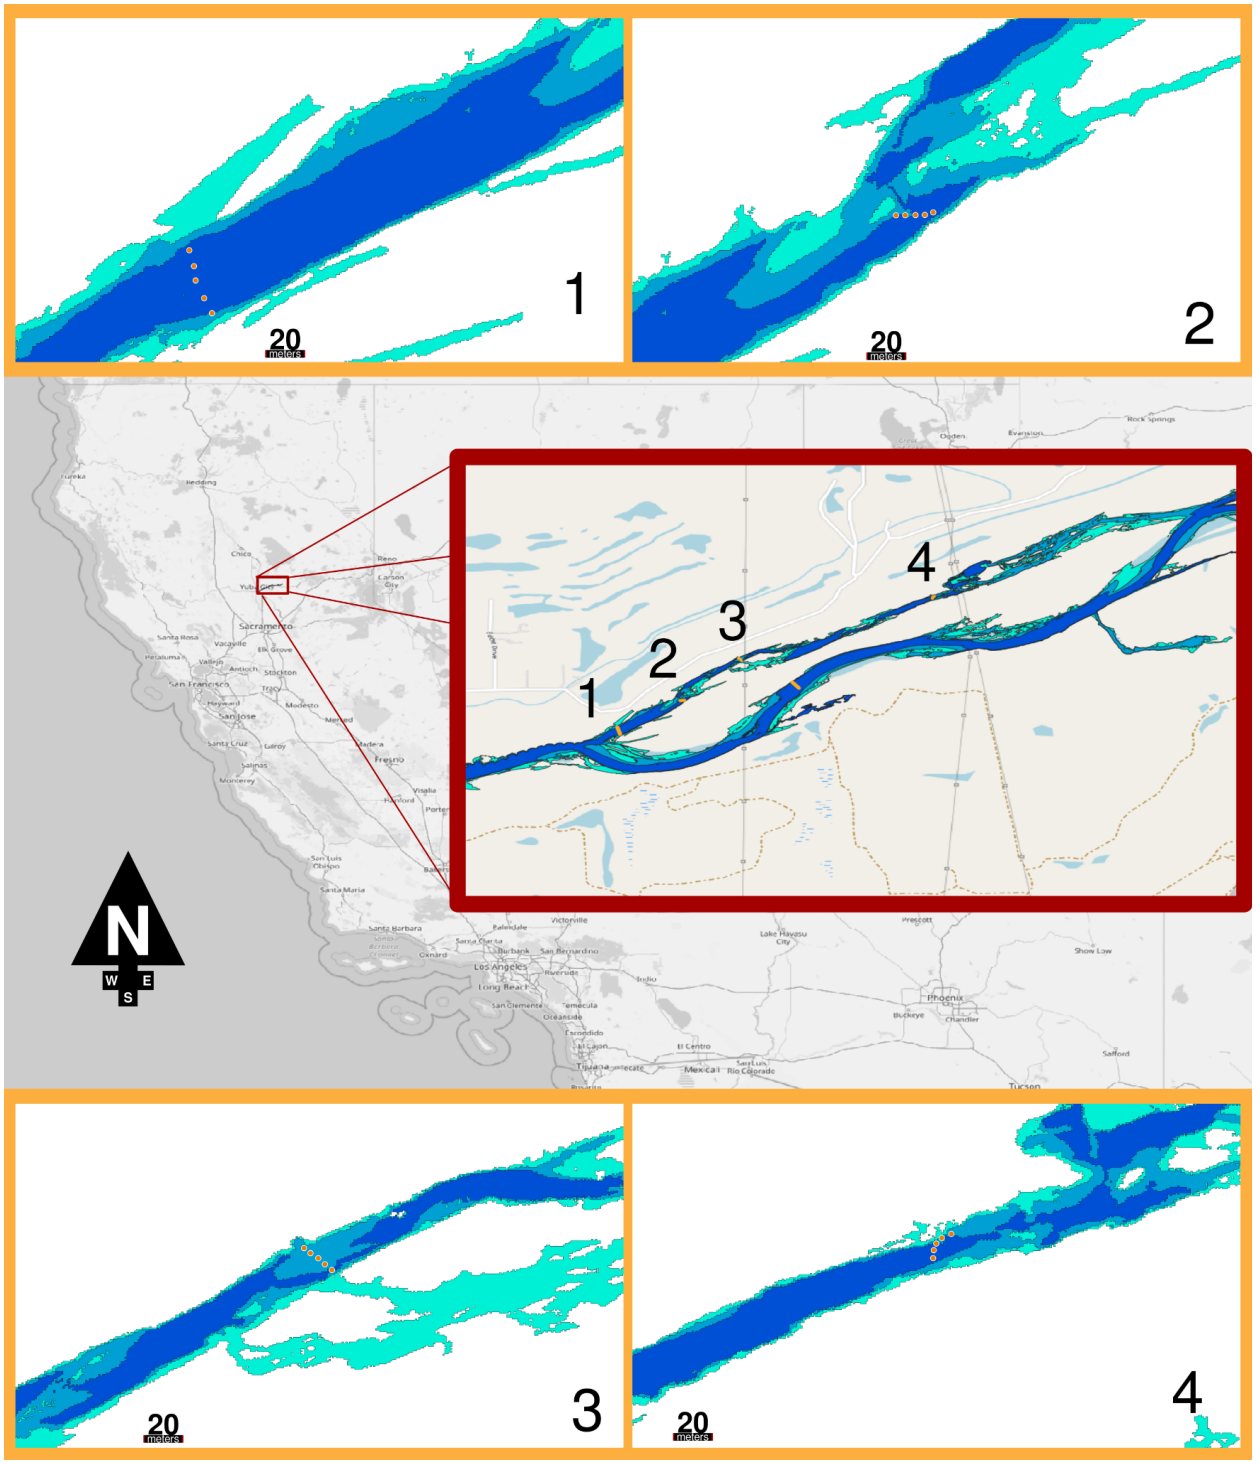

Figure 1. Map of the study reach in the lower Yuba River. At top-right of inset map (red), is Daguerre Point Dam. The river flows from the top right to the bottom left. Orange lines correspond to sensors (1-4), which span the channel at that point. Orange dots represent the location of the individual sensors. The darkest blue areas are wetted at a flow of 8.5 m^3/s, the medium or turquoise are wetted at a flow of 28 m^3/s, and the lightest teal wetted at a flow of 141 m^3/s

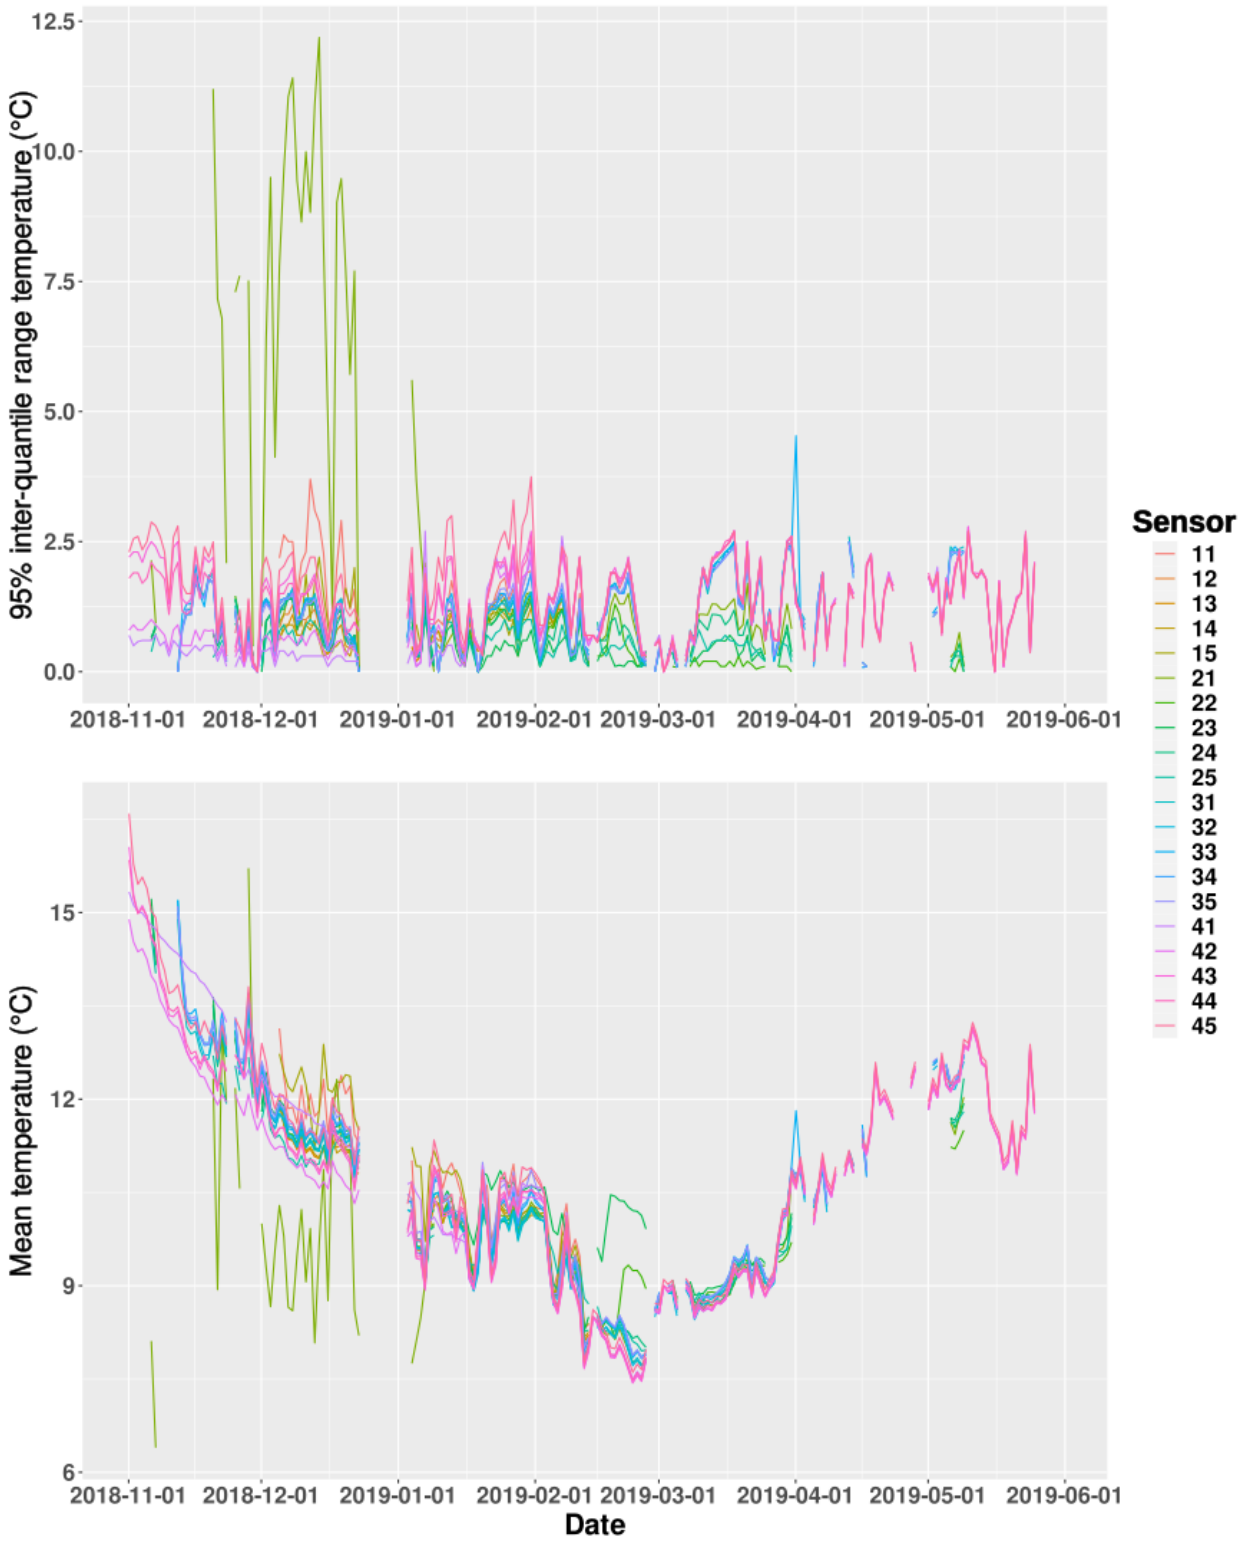

Figure 2. Daily (a) 95% inter-quantile range and (b) mean observed temperatures for each of the 20 sensors installed in the sidechannel of the study reach. The first number of each sensor name corresponds to the sensor node, the second corresponds to the sensor's position within the string of temperature sensors.

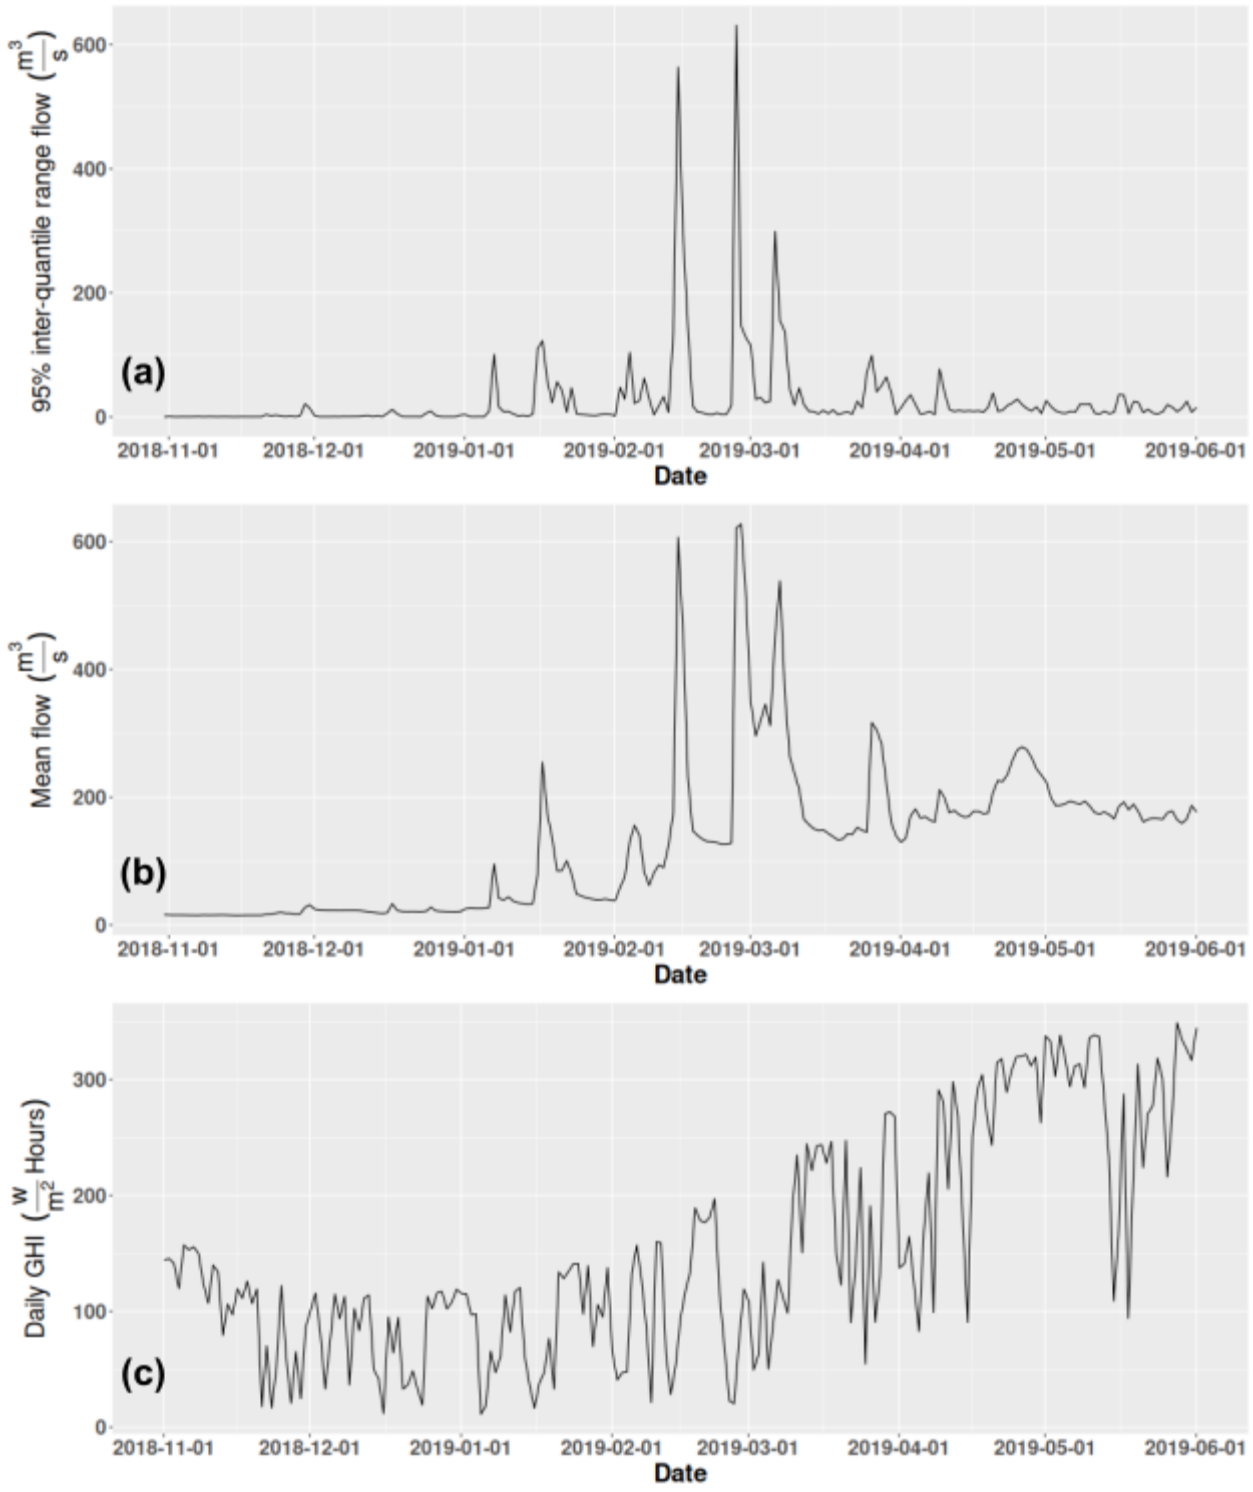

Figure 3. Environmental data from sensors near the study reach over the period of sensor deployment, including (a) daily 95% inter-quantile range and (b) mean for flows, as measured by the Marysville USGS (2022) gauge; and (c) sum of hourly solar radiation (GHI) from Beale Air Force Base (NSRDB 2022).

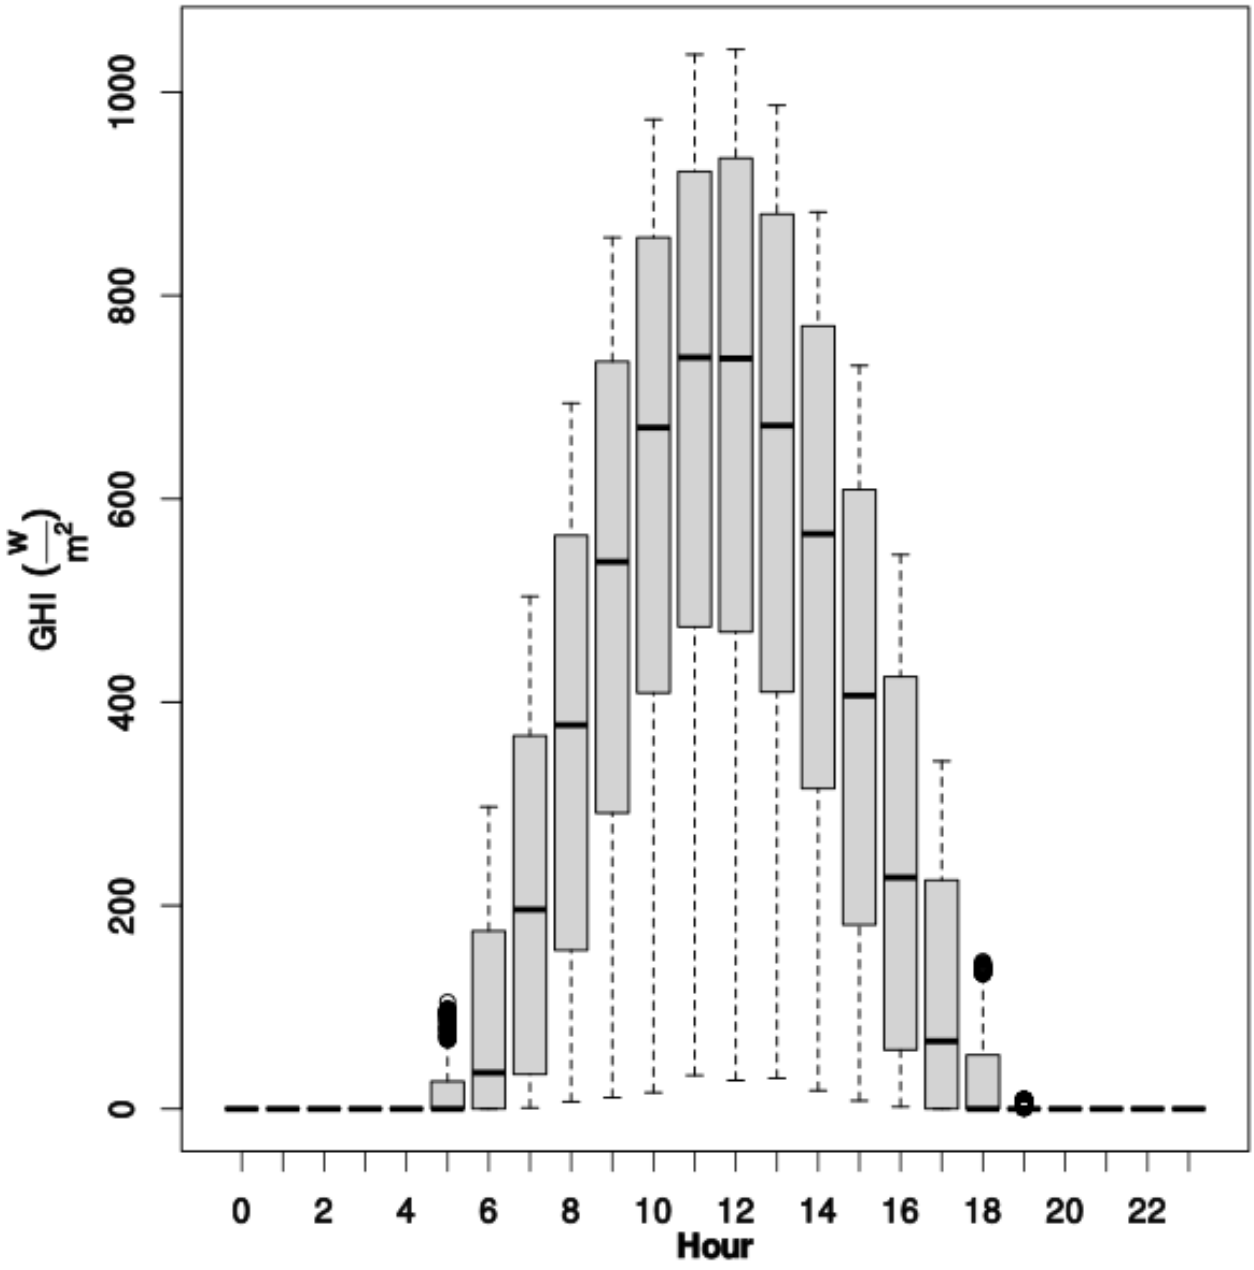

Figure 4. Boxplot of solar radiation (GHI; NSRDB 2022) data from 01-Nov-2018 through 01-Jun-2019, grouped by hour of day (PST).

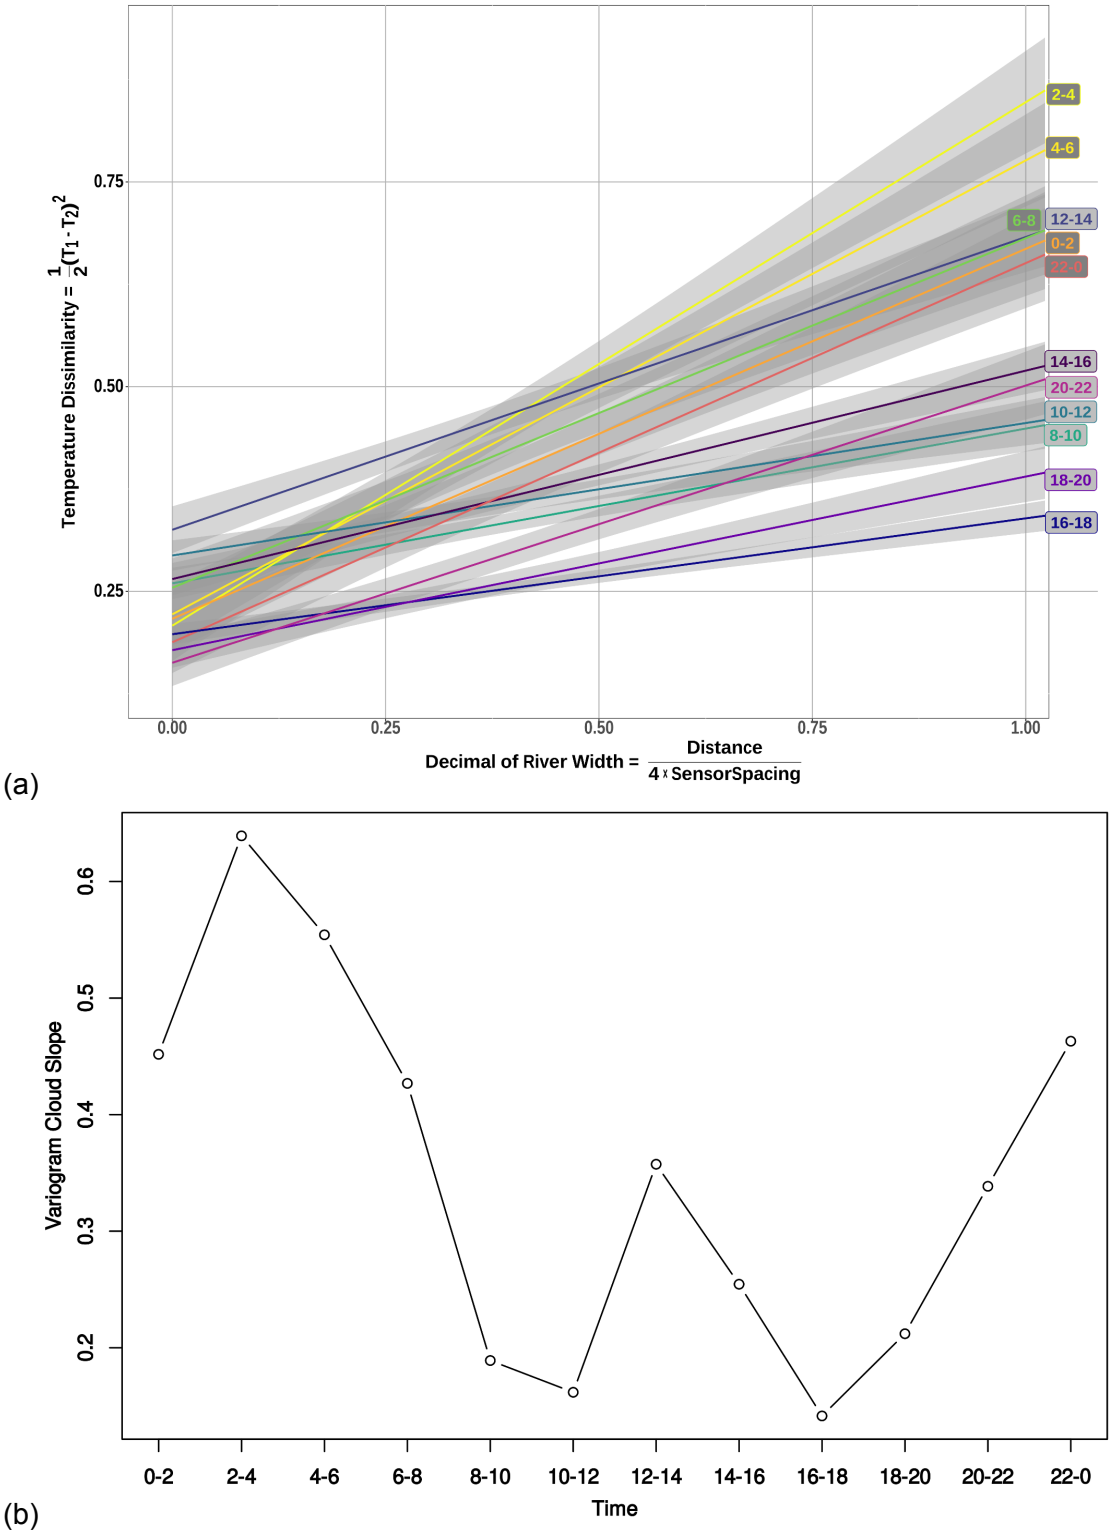

Figure 5. (a) Lateral semi-variance relative to percent of channel-width moved, broken out by 2-hour time-of-day block. Curves were generated via a semi-variogram cloud, which was then smoothed using a linear model. Shaded regions are 95% CI. (b) Slopes of smoothed semi-variance curves for each time bucket – positive slopes reflect positive spatial autocorrelation. Slopes near zero are not autocorrelated.

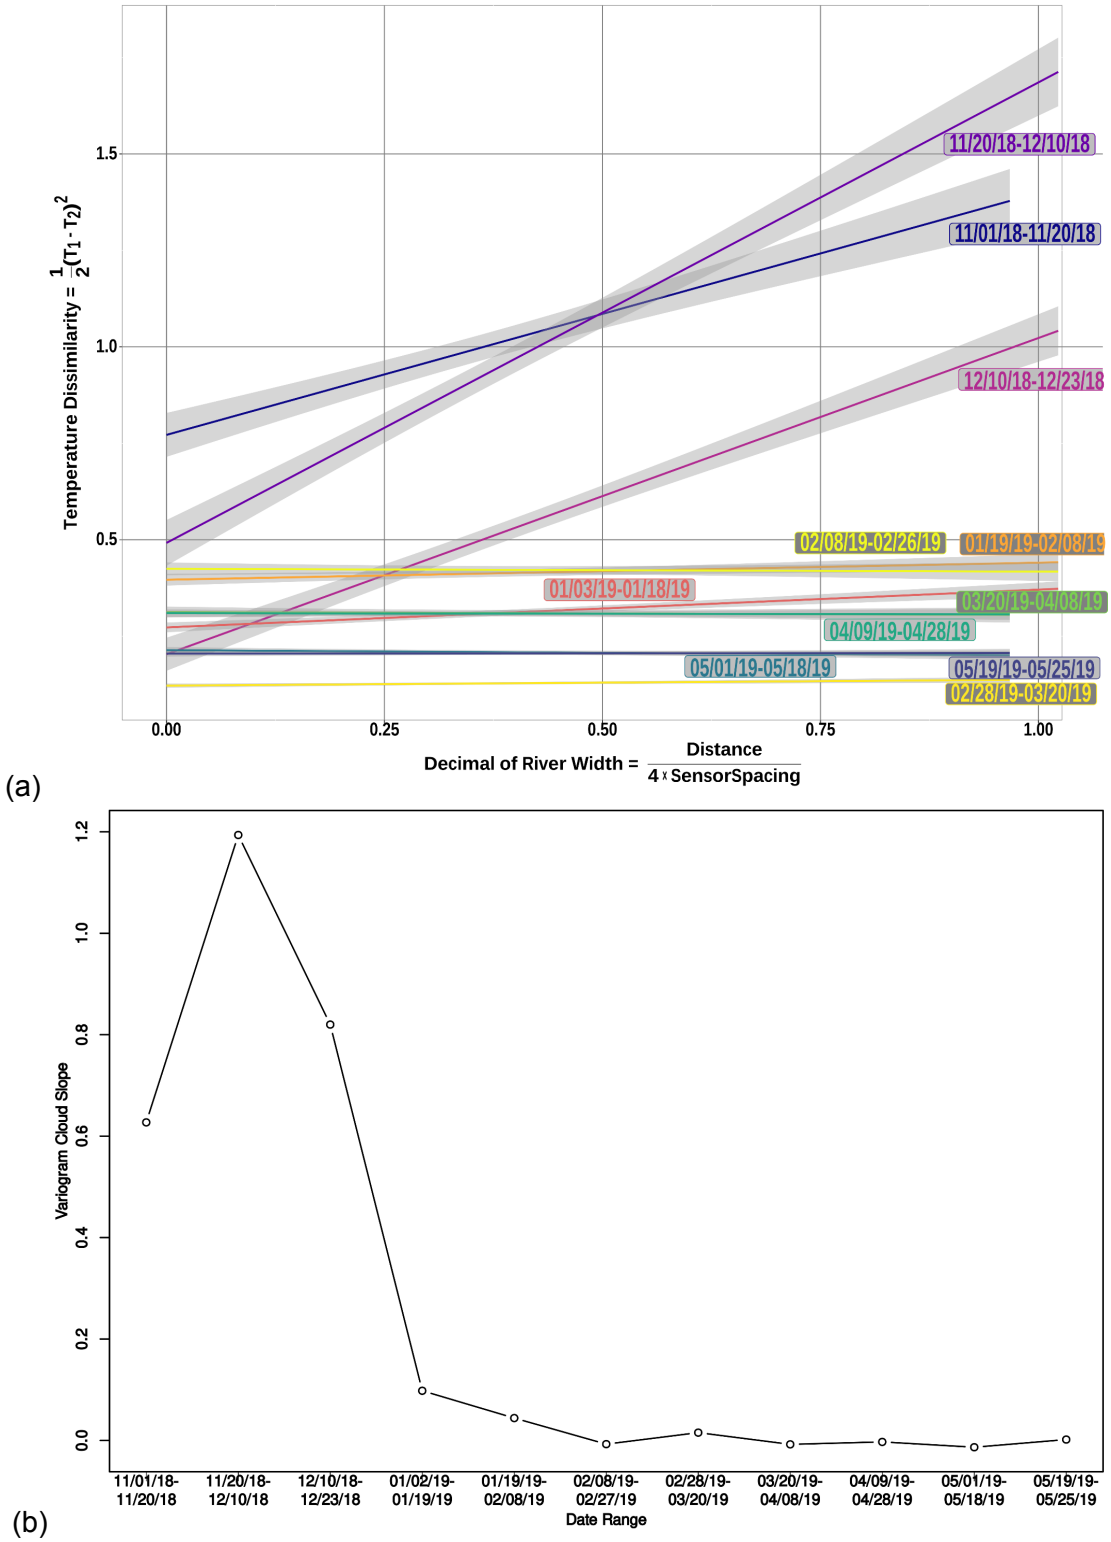

Figure 6. (a) Lateral semi-variance relative to percent of channel-width moved, broken out by 20-day sampling block. Curves were generated via a semi-variogram cloud, which was then smoothed using a linear model. Shaded regions are 95% CI. (b) Slopes of smoothed semi-variance curves for each time bucket – positive slopes reflect positive spatial autocorrelation. Slopes near zero are not autocorrelated.

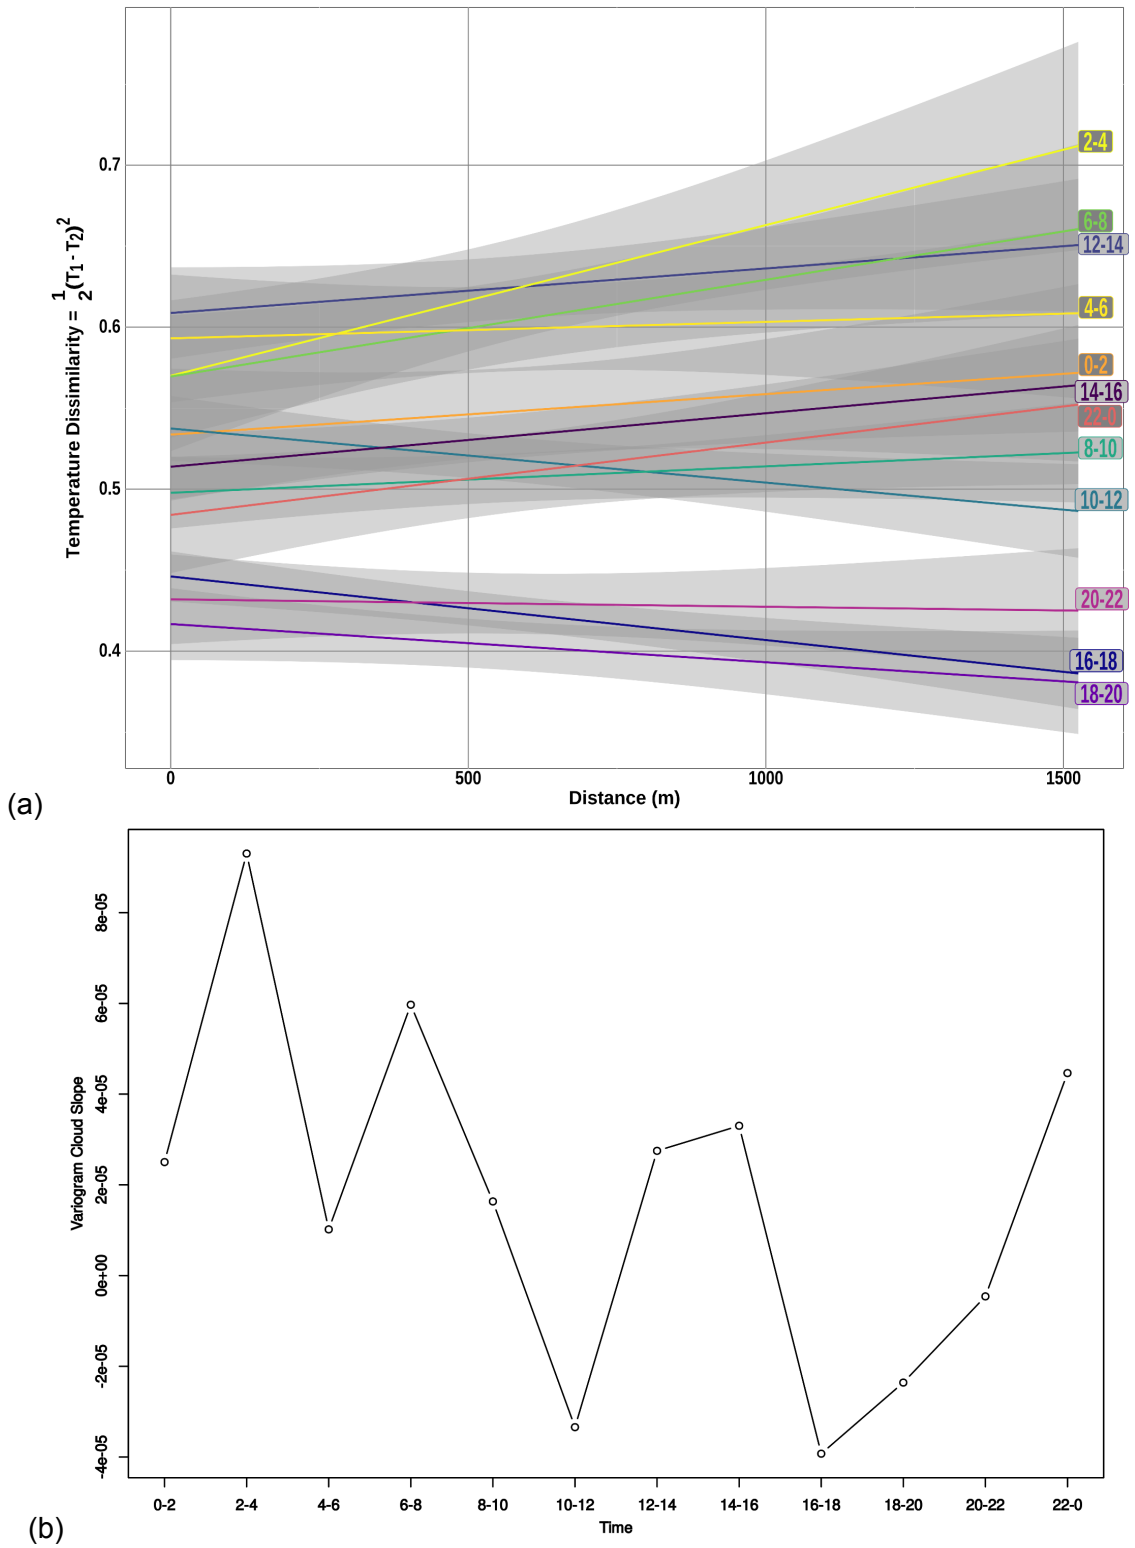

Figure 7. (a) Longitudinal semi-variance relative to stream-oriented distance moved, broken out by two-hour time-of-day block. Curves were generated via a semi-variogram cloud, which was then smoothed using a linear model. Shaded regions are 95% CI. (b) Slopes of smoothed semi-variance curves for each time bucket – positive slopes reflect positive spatial autocorrelation. Slopes near zero are not autocorrelated.

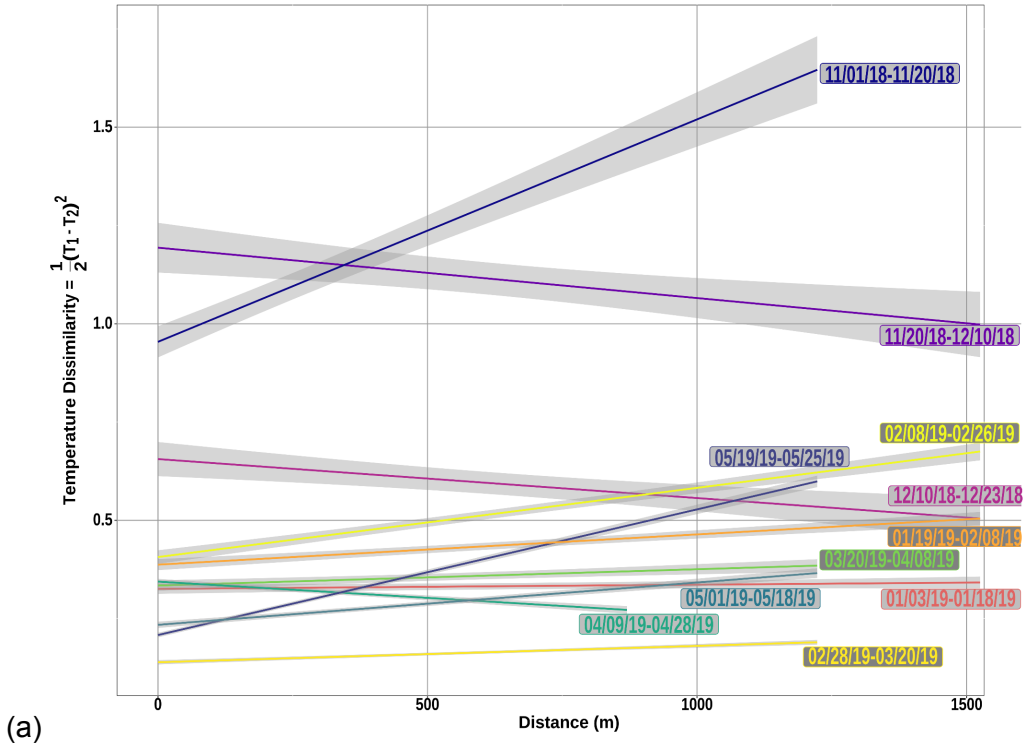

(a)

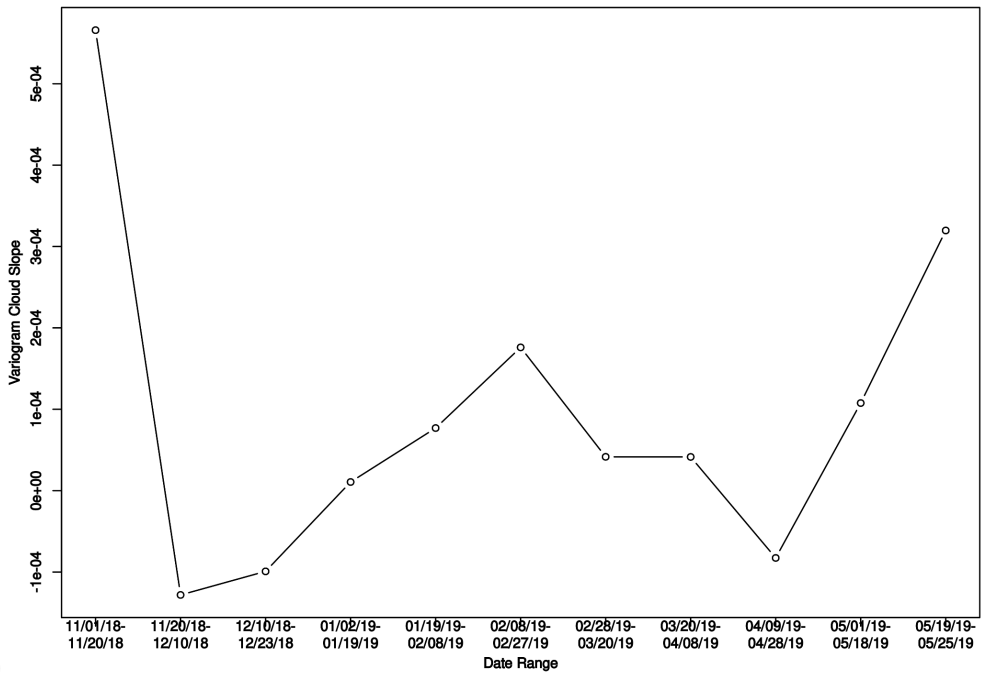

(b)

Figure 8. (a) Longitudinal semi-variance relative to stream-oriented distance moved, broken out by 20-day sampling block. Curves were generated via a semi-variogram cloud, which was then smoothed using a linear model. Shaded regions are 95% CI. (b) Slopes of smoothed semi-variance curves for each time bucket – positive slopes reflect positive spatial autocorrelation. Slopes near zero are not autocorrelated.