Space-time large deviations in capacity-constrained relay networks

Christian Hirsch, Benedikt Jahnel and Robert Patterson

Department of Mathematical Sciences, Aalborg University,
Skjernvej 4,
9220 Aalborg Ø, Denmark.
E-mail address: christian@math.aau.dk
URL: http://people.math.aau.dk/~christian/index.html

Weierstrass Institute for Applied Analysis and Stochastics,
Mohrenstraße 39,
10117 Berlin, Germany.
E-mail address: benedikt.jahnel@wias-berlin.de
URL: www.wias-berlin.de/people/jahnel

Weierstrass Institute for Applied Analysis and Stochastics,
Mohrenstraße 39,
10117 Berlin, Germany.
E-mail address: robert.patterson@wias-berlin.de
URL: www.wias-berlin.de/people/patterso

Abstract. We consider a single-cell network of random transmitters and fixed relays in a bounded domain of Euclidean space. The transmitters arrive over time and select one relay according to a spatially inhomogeneous preference kernel. Once a transmitter is connected to a relay, the connection remains and the relay is occupied. If an occupied relay is selected by another transmitters with later arrival time, this transmitter becomes frustrated. We derive a large deviation principle for the space-time evolution of frustrated transmitters in the high-density regime.

1. Introduction and main results

We consider a single-cell communication network of random *transmitters* and fixed *relays*. Every transmitter tries to send data to a central entity via one relay according to a spatially dependent preference function. Each relay can only serve

one transmitter and the transmitters are competing for this shared capacity. In particular, a group of transmitters might not be successful in finding relays to release their data to and therefore become *frustrated*. We will assume that the start of the data transmission is time dependent so that the set of frustrated transmitters gradually increases over time. We present a large deviation principle (LDP) for the measure-valued process of frustrated transmitters when their number increases.

The motivation for this work is to derive the LDP for a capacity-constrained network embedded in the Euclidean space. So far in the literature, there have been two separate approaches. On the one hand, considerable work has been done to understand the large deviation behavior of sophisticated capacity-constrained networks in a mean-field setting, see for example Graham and Méléard (1997b,a). On the other hand, driven by recent developments in wireless networks, from an engineering perspective, there has been a surge in research activities to develop a fundamental understanding of spatial effects in models that are based on stochastic geometry Bethanabhotla et al. (2016); Sankararaman and Baccelli (2017). In the present paper, we analyze a simple model of a spatial relay network which is to be seen as a first step into the realm of space-time LDPs for capacity-constrained networks.

More specifically, consider fixed relays at locations $Y^\lambda = (y_j)_{1 \le j \le n_\lambda}$ in a compact window $W \subset \mathbb{R}^d$ with boundaries of vanishing Lebesgue measure. We investigate the high-intensity regime and thus assume that the empirical distribution

$$l_\lambda = \lambda^{-1} \sum_{y_j \in Y^\lambda} \delta_{y_j}$$

converges weakly, as $\lambda \uparrow \infty$, to some probability measure μ_R on W. Further there will be transmitters distributed according to a Poisson point process X^λ in W. Its intensity measure is of the form $\lambda \mu_\mathsf{T}^\mathsf{s}$ with $\lambda > 0$ and $\mu_\mathsf{T}^\mathsf{s}$ a finite measure on W which is absolutely continuous w.r.t. the Lebesgue measure. Initially, all relays are idle and the transmitters do not send data. For each transmitter X_i there is a randomly distributed time $T_i \in [0, t_\mathsf{f}]$ and in (T_i, t_f) there will be constant data transmission. The times T_i are assumed to be iid with distribution $\mu_\mathsf{T}^\mathsf{t}$ which is absolutely continuous w.r.t. the Lebesgue measure on $[0, t_\mathsf{f}]$.

The transmitters are assumed to have basic knowledge about the transmission quality to each of the relays. More precisely, each transmitter X_i uses this information to assign a spatial preference $\kappa(X_i, y_j) \in [0, 1]$ for the connection from X_i to y_j. At time T_i the user X_i tries to send its data to a relay y_j chosen according to the *preference kernel*

$$\kappa(y_j | X_i) = \frac{\kappa(X_i, y_j)}{\sum_{y_k \in Y} \kappa(X_i, y_k)}, \tag{1.1}$$

so that the selection probability is proportional to the spatial preference function κ. In our model, data transmission fails if the chosen relay is already occupied by some other transmitter that has established a connection earlier in time.

Since every transmitter X_i keeps the connection until time t_f, the time-dependent status of its target relay $Y^{(i)}(t) \in \{0, 1\}$ is a step function, starting in 0 as being idle and jumping to 1 at the time where it becomes busy. In particular $\Gamma^\lambda = \{\Gamma_t^\lambda\}_{0 \le t \le t_\mathsf{f}}$

with

$$\Gamma_t^\lambda = \frac{1}{\lambda} \sum_{X_i \in X} \mathbb{1}\{X_i(t) = 1\}\delta_{X_i} \tag{1.2}$$

denotes the normalized, time-dependent random measure of frustrated transmitters. Here X_i is the (time invariant) position of a relay, but $X_i(t)$ represents the time-dependent status of frustration of the transmitter X_i. More precisely, $X_i(t)$ equals 0 for $t < T_i$ and jumps to 1 at $t = T_i$ if the chosen relay is already occupied, i.e. $Y^{(i)}(t-) = 1$. Figure 1.1 provides a snapshot of the relay network after a finite time.
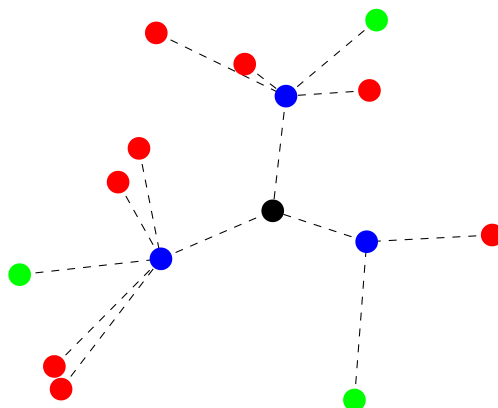

FIGURE 1.1. Collection of transmitters (green and red) communicating with one central entity (black) via one relay (blue) each. Red transmitters are frustrated due to low capacity at the associated relay. Green transmitters are satisfied as they successfully send data to a relay.

From the perspective of the network operator, it is critical to understand the time- and space-dependent process of frustrated transmitters in large networks. Once this understanding is achieved, it is possible to answer questions like:

What is the overall proportion of transmitters which are frustrated at a given point in time?   Are most of the frustrated transmitters located in a specific area?

Our main result provides a probabilistic description of the process of frustrated transmitters in an asymptotic regime where the number of devices tends to infinity. In particular, we perform a large deviation analysis for the empirical measure of frustrated transmitters Γ^λ, where rates of convergence for unlikely events are derived.

1.1. *Large deviations without preference.* Note that Γ^λ allows us to keep track of transmitter locations but not on the chosen connections to specific relays. In particular, for general preference kernels, Γ^λ is not Markovian since the required spatial information of occupied relays for a new transmitter request at time t cannot be extracted from Γ_{t-}^λ. However, for $\kappa \equiv 1$, the Markovianity of Γ^λ can be preserved since transmitters have no spatial preference in their choice of relays. Therefore, we establish the case $\kappa \equiv 1$ first and use it as a basis for the general case.

For this, we first note that the process of frustrated transmitters Γ^λ can be recovered from the process of *satisfied transmitters* $B^\lambda = \{B_t^\lambda\}_{0 \le t \le t_f}$. This process is easier to describe as it coincides with the number of busy relays. More precisely, when at time $t = T_i$ a transmitter request from X_i arrives, the chosen relay is already busy with probability given by the total proportion of idle relays and the transmitter becomes frustrated. In this case, B^λ stays constant at t. Otherwise, if the chosen relay is idle, the relay becomes busy and B^λ grows by $\lambda^{-1}\delta_{X_i}$ at t. As the number of satisfied transmitters equals the number of busy relays, this has probability $1 - B_{t-}^\lambda(W)/r_\lambda$ where $r_\lambda = |Y^\lambda|/\lambda$. Note that this random choice of relays can be encoded by assigning a uniform random variable $U_i \in [0,1]$ to transmitter X_i. If $U_i \in [B_{t-}^\lambda(W)/r_\lambda, 1]$ then X_i connects to an idle relay. The encoding of the spatial relay configuration into a $[0,1]$-valued threshold is illustrated in Figure 1.2.
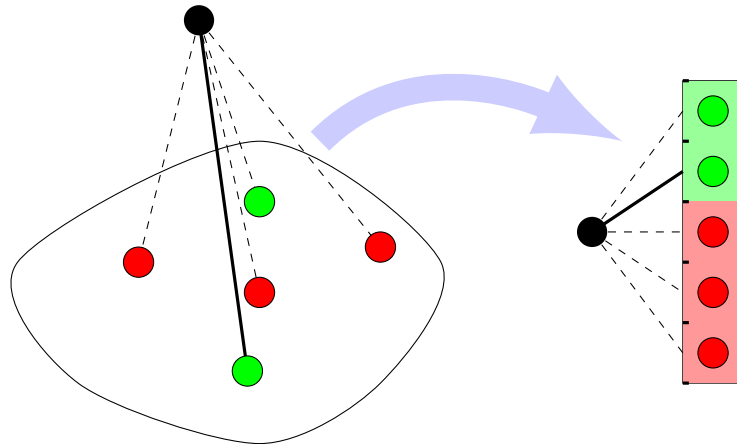

FIGURE 1.2. A transmitter (black) chooses a relay at random (left), where the relay can either be already busy (red) or idle (green). Without spatial preferences the relay information can be reduced to a single threshold in $[0,1]$ (right).

More precisely, the evolution of the process of satisfied transmitters B^λ is given as the solution of the time integral equation

$$B_t(\mathrm{d}x) = \int_0^t \int_{B_{s-}(W)/r_\lambda}^1 L_\lambda(\mathrm{d}s, \mathrm{d}u, \mathrm{d}x) = \int_0^t L_\lambda(\mathrm{d}s, [B_{s-}(W)/r_\lambda, 1], \mathrm{d}x), \quad (1.3)$$

where

$$L_\lambda = \lambda^{-1} \sum_{X_i \in X^\lambda} \delta_{(T_i, U_i, X_i)}$$

is the empirical measure of transmitters. In particular, the Poisson point process X^λ of transmitters carries marks for data request time and choice variable. Its intensity measure $\lambda\mu_T$ in $V = [0, t_f] \times [0,1] \times W$ is given by

$$\mu_T = \mu_T^t \otimes U \otimes \mu_T^s$$

and U is the uniform distribution on $[0,1]$.

In words, equation (1.3) describes the evolution of the empirical measure of satisfied transmitters. This measure gains mass λ^{-1} at position x if there is an additional transmitter request from position x at time t and this transmitter picks a relay that is idle at time $t-$. From B^λ the random measure of frustrated transmitters can be recovered via

$$\Gamma_t^\lambda(\mathrm{d}x) = L_\lambda([0,t],[0,1],\mathrm{d}x) - B_t^\lambda(\mathrm{d}x).$$

Note that, as the number of devices tends to infinity, in the limit, the point masses associated to individual devices disappear and the picture becomes continuous. Thus, we need to introduce processes of frustrated transmitters for absolutely continuous measures. More precisely, let ν be a finite measure in $\mathcal{M} = \mathcal{M}(V)$ where $\nu(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x)$ has the interpretation of the intensity of transmitters at $\mathrm{d}x$ with data-entry time $\mathrm{d}t$ and choice variable $\mathrm{d}u$. Then, for general measures and normalized number of relays r, equation (1.3) has the form

$$\beta_t(\mathrm{d}x) = \int_0^t \int_{\beta_{s-}(W)/r}^1 \nu(\mathrm{d}s, \mathrm{d}u, \mathrm{d}x) = \int_0^t \nu(\mathrm{d}s, [\beta_{s-}(W)/r, 1], \mathrm{d}x). \qquad (1.4)$$

For illustration purposes let us present two examples here.

(1) For the empirical measure $\nu = L_\lambda$ of the Poisson point processes X^λ and $r = r_\lambda$, the unique solution β for (1.4) is given by B^λ.
(2) For the a priori measure μ_T as a driver and normalized relay number $r = 1$, the unique solution of (1.5) is given by $\beta\mu_\mathsf{T}^\mathsf{s}$ where

$$\beta_t = \mu_\mathsf{T}^\mathsf{s}(W)^{-1}(1 - e^{-\mu_\mathsf{T}^\mathsf{s}(W)\mu_\mathsf{T}^\mathsf{t}([0,t])}).$$

Note that, instead of (1.4) with $r = 1$, it suffices to consider the scalar equation

$$\beta_t = \int_0^t \nu(\mathrm{d}s, [\beta_{s-}, 1], W). \qquad (1.5)$$

Using Schauder's fixed point theorem and monotonicity, we will show in Subsection 2.2 that existence and uniqueness of solutions of (1.5) for absolutely continuous measures can be established. With the steps we have just described we arrive at a solution which we will denote $\beta(\nu)$. Moreover, in the absolutely-continuous case, the solution will be continuous and increasing in time.

From the solution $\beta(\nu)$ one can compute the normalized process of frustrated transmitters $\gamma(\nu)$ via the formula

$$\gamma_t(\nu)(\mathrm{d}x) = \nu([0,t],[0,1],\mathrm{d}x) - \beta_t(\nu)(\mathrm{d}x). \qquad (1.6)$$

In the following theorem we show the LDP for Γ^λ in the setting where $\kappa \equiv 1$. Recall the definition of the relative entropy

$$h(\nu|\mu) = \int \log \frac{\mathrm{d}\nu}{\mathrm{d}\mu}\mathrm{d}\nu - \nu(V) + \mu(V)$$

if $\nu \ll \mu$ with $h(\nu|\mu) = \infty$ otherwise. We consider Γ^λ as a measure-valued process and work in the Skorohod space. That is, we consider

$$\mathbb{D} = \{f \in \mathcal{M}(W)^{[0,t_f]} : f \text{ is càdlàg w.r.t. the weak topology on } \mathcal{M}(W)\}$$

equipped with the Skorohod topology, for details see for example Feng and Kurtz (2006).

Theorem 1.1. *The family of measure-valued processes* Γ^λ *satisfies the LDP in* $\mathbb{D}$ *with good rate function given by* $I(\gamma) = \inf_{\nu \in \mathcal{M}: \gamma(\nu) = \gamma} h(\nu | \mu_\mathsf{T})$.

Note that a scalar variant of Theorem 1.1 appears in Boucheron et al. (2002, Theorem 2.7) when transmitters are interpreted as bins and relays as balls. This provides an application of this classical model from random discrete structures to communication networks. However, the results in Boucheron et al. (2002) cannot be used to prove Theorem 1.1, since the balls arrive at deterministic times, whereas in our setting also the arrival times of the transmitters are random and not necessarily homogeneous.

Moreover, in the scalar setting, it is possible to explicitly perform the minimization only over the choice component. Then, the rate function of Theorem 1.1 can be expressed as

$$I(\gamma) = \inf_\beta \int_0^{t_\mathsf{f}} \big[h(\dot{\beta}_t | 1 - \beta_t) + h(\dot{\gamma}_t | \beta_t) \big] \mathrm{d}t.$$

Here, we assumed $\mu_\mathsf{T}^\mathsf{t}(\mathrm{d}t) = \mathbb{1}_{[0,t_\mathsf{f}]}(t)\mathrm{d}t$ for simplicity and the infimum is taken over increasing and absolutely continuous paths. This form of the rate function coincides with the one derived in Shwartz and Weiss (1995) in the setting of chemical-reaction networks. The interpretation of our communication-network evolution is then that the species of frustrated and satisfied transmitters are generated at rates β_t respectively $1 - \beta_t$. Theorem 1.1 is not covered by the standard results in Léonard (1995); Shwartz and Weiss (1995), since in our setting, these rates are not bounded away from zero.

In the literature attempts have been made to relax the assumption of strictly positive rates (Shwartz and Weiss, 2005). However, our model is also not covered by the results in Shwartz and Weiss (2005), as the crucial *interior cone property* is violated: Once all relays are occupied it is not possible to move back to a state where satisfied users are generated at positive rate.

Finally, it is difficult to extend the interpretation of our network as a chemical reaction if we take spatial resolution into account. Then, the space of species would become uncountable since spatial locations have to be tracked. This is another reason for our decision to rely on marked Poisson point processes and measure-valued differential equations in Theorem 1.1.

1.2. *Large deviations with preference.* In this section, we deal with spatial preferences of transmitters. As a consequence, the probability to send to a certain location depends on the spatial location of transmitters and relays and not just on the number of relays in the entire domain. In particular, the encoding into a single $[0, 1]$-valued threshold falls short of capturing the information required for describing the evolution of frustrated transmitters. However, for a sufficiently smooth preference kernel, for a given transmitter, the relay choice is approximately uniform in a neighborhood around any given relay location. Therefore, as an approximation we partition W into a finite number of patches. Then to each of these patches we associate a separate $[0, 1]$-valued threshold describing the approximate proportion of busy relays in that patch. This encoding is illustrated in Figure 1.3.

Here, we assume μ_R to be absolutely continuous w.r.t. the Lebesgue measure. The spatial preference function may be discontinuous and non-positive, but for sufficiently high densities from any transmitter location it should be possible to
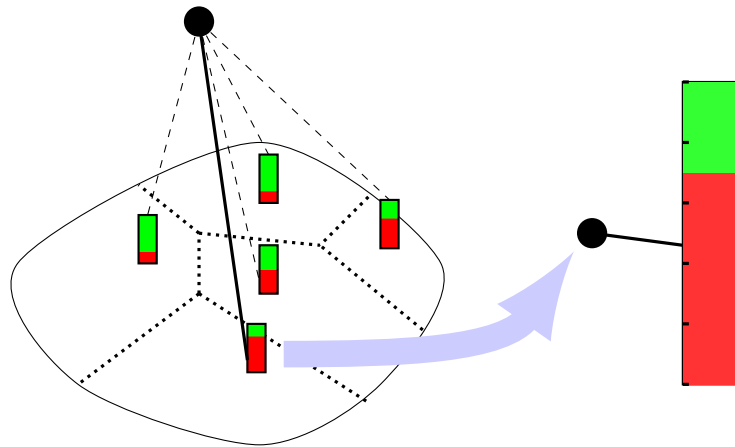
FIGURE 1.3. A transmitter (black) chooses a relay location at a coarse scale according to a spatial preference function (left). At a fine scale the configuration of busy relays can be encoded in a $[0, 1]$-valued threshold as in the setting of flat preference kernels (right).

connect to some relay location. More precisely, we assume that κ is jointly continuous $\mu_\mathsf{T}^\mathsf{s} \otimes \mu_\mathsf{R}$-almost everywhere and for all $x \in W$ there exists $y \in W$ such that $\kappa(x, y) > 0$, $y \in \operatorname{supp}(\mu_\mathsf{R})$ and (x, y) is a continuity point of κ. Recall that a transmitter at location $x \in W$ chooses a relay at location $\mathrm{d}y$ with probability

$$\kappa_{l_\lambda}(\mathrm{d}y | x) = \frac{\kappa(x, y)}{\int \kappa(x, z) l_\lambda(\mathrm{d}z)} l_\lambda(\mathrm{d}y) \tag{1.7}$$

where, by our assumption, for sufficiently large λ the denominator is bounded away from 0 uniformly in x.

Next we consider the interplay between the spatial location of the relays and how the spatial preferences of the transmitters evolve in time. Note that the process of transmitter requests to a relay at location $\mathrm{d}y$ is a Poisson point process Z^λ on $V' = V \times W$ with intensity measure

$$\mu(l_\lambda)(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, \mathrm{d}y) = \kappa_{l_\lambda}(\mathrm{d}y | x) \mu_\mathsf{T}(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x).$$

As in Theorem 1.1 the LDP will be obtained by making use of the observation that the measure of frustrated transmitters Γ^λ can be interpreted as a functional of the empirical measure $\mathfrak{L}_\lambda \in \mathcal{M}' = \mathcal{M}(V')$ associated to Z^λ. In particular, in the large deviations it is possible for the process to distort the new a priori measure $\mu(\mu_\mathsf{R})$ into an absolutely-continuous transmitter-request distribution $\mathfrak{n}$. More precisely, we generalize (1.5) into the measure-valued time integral equation

$$\mathbf{b}_t(\mathrm{d}x, \mathrm{d}y) = \int_0^t \int_{\frac{\mathrm{d}\mathbf{b}_{s-}(W, \cdot)}{\mathrm{d}\nu_\mathsf{R}}(y)}^1 \mathfrak{n}(\mathrm{d}s, \mathrm{d}u, \mathrm{d}x, \mathrm{d}y) = \int_0^t \mathfrak{n}(\mathrm{d}s, [\tfrac{\mathrm{d}\mathbf{b}_{s-}(W, \cdot)}{\mathrm{d}\nu_\mathsf{R}}(y), 1], \mathrm{d}x, \mathrm{d}y) \tag{1.8}$$

where we allow the relay measure ν_R also to be general in order to cover both, empirical measures as well as absolutely-continuous measures. For example,

(1) if the driving measure is given by $(\mathfrak{L}_\lambda, l_\lambda)$, then (1.8) has a unique solution $\mathbf{b}(\mathrm{d}x, \mathrm{d}y)$ and, for $\kappa \equiv 1$, $B^\lambda(\mathrm{d}x) = \mathbf{b}(\mathrm{d}x, W)$ in distribution.

(2) if $\mathfrak{n} = \mu(\mu_\mathsf{R})$, then a solution is given by

$$\mathbf{b}_t(\mathrm{d}x, \mathrm{d}y) = (1 - e^{-\mu_\mathsf{T}^\mathsf{t}([0,t]) \int \kappa(y|z)\mu_\mathsf{T}^\mathsf{s}(\mathrm{d}z)}) \frac{\kappa(y|x)\mu_\mathsf{T}^\mathsf{s}(\mathrm{d}x)}{\int \kappa(y|z)\mu_\mathsf{T}^\mathsf{s}(\mathrm{d}z)} \mu_\mathsf{R}(\mathrm{d}y).$$

Note that, for general driving measures $(\mathfrak{n}, \nu_\mathsf{R})$, existence of solutions for equation (1.8) is unclear. However, existence of solutions can be established if we assume $\mathfrak{n}$ to be of the form

$$\mathfrak{n}(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, \mathrm{d}y) = \mathfrak{n}_y(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x)\nu_\mathsf{R}(\mathrm{d}y)$$

where $(\mathfrak{n}_y)_{y \in W}$ is a transition kernel $W \to \mathcal{M}$ such that every $\mathfrak{n}_y(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x)$ is absolutely continuous w.r.t. μ_T. Indeed, then a solution of (1.8) is given by

$$\mathbf{b}(\mathrm{d}x, \mathrm{d}y) = \beta(\mathfrak{n}_y)(\mathrm{d}x)\nu_\mathsf{R}(\mathrm{d}y) \tag{1.9}$$

and will be denoted by $\mathbf{b}(\mathfrak{n}, \nu_\mathsf{R})$.

If $\mathbf{b}(\mathfrak{n}, \nu_\mathsf{R})$ is well-defined, we can compute the process of frustrated transmitters $\gamma(\mathfrak{n}, \nu_\mathsf{R})$ via the formula

$$\gamma_t(\mathfrak{n}, \nu_\mathsf{R})(\mathrm{d}x) = \mathfrak{n}([0,t], [0,1], \mathrm{d}x, W) - \mathbf{b}_t(\mathfrak{n}, \nu_\mathsf{R})(\mathrm{d}x, W)$$

and in particular $\Gamma^\lambda = \gamma(\mathfrak{L}_\lambda, l_\lambda)$ equals (1.2) in distribution. Now, we present the main result of this paper.

Theorem 1.2. *The family of measure-valued processes Γ^λ satisfies the LDP in $\mathbb{D}$ with good rate function given by $I(\gamma) = \inf_{\mathfrak{n} \in \mathcal{M}': \gamma(\mathfrak{n}, \mu_\mathsf{R}) = \gamma} h(\mathfrak{n}|\mu(\mu_\mathsf{R}))$.*

Note that at first sight, Γ^λ is an object that requires knowledge on how many transmitters choose a relay and not just the number of transmitters in a given area that choose relays in another given area. Therefore, it may come as a surprise that we can state Theorem 1.2 as a measure-valued LDP and not as a LDP on the level of spatial configurations. To reconcile Theorem 1.2 with the reader's intuition, we note that after approximation by flat preference functions, we deal with an independent collection of processes of the type considered in Theorem 1.1. This allows us to aggregate the information about an entire local configuration of occupied relays into a single number.

The rest of the manuscript is organized as follows. Sections 2 and 4 provide high-level overviews for the proofs of Theorems 1.1 and 1.2, respectively. Detailed proofs for all supporting results can be found in Sections 3 and 5.

2. Proof of Theorem 1.1

The idea for the proof of Theorem 1.1 is to use a Sanov-type result for the transmitter distribution ν and apply the contraction principle for solutions of the associated differential equation (1.5). The main ingredient in this approach is then the continuity of solutions at user distributions with finite entropy. The theory of ODE provides us with conditions under which continuity of solutions w.r.t. parameters can be inferred. Unfortunately, these results mostly work under Lipschitz assumptions which are stronger than the finite entropy bounds provided in our setting. However, we construct a two-step Picard approximation that is tailor-made to provide the right balance between two opposing constraints: It is simple enough to

be continuous with respect to the driving measure and at the same sufficiently close to the true solution to satisfy uniform approximation properties. This allows us to employ the LDP tool of exponential approximations Dembo and Zeitouni (1998, Definition 4.2.14) to complete the proof of Theorem 1.1.

2.1. *Markovian structure of* Γ^λ. In this subsection we show that the solutions to equation (1.3), modeled using the random variable U_j, is equal in distribution to the process of frustrated users as defined in (1.2). To make this precise, let $\tilde{\Gamma}^\lambda$ denote the unique solution of (1.3) for the initial condition $\tilde{\Gamma}_0^\lambda = 0$.

Proposition 2.1. *The processes* Γ^λ *and* $\tilde{\Gamma}^\lambda$ *have the same distribution.*

Proof: Let us fix a realization of $N^\lambda = |X^\lambda|$ many transmitter locations $(x_i)_{1 \leq i \leq N^\lambda}$ together with their data transmission times $(t_i)_{1 \leq i \leq N^\lambda}$. After that, the measure-valued process $\Gamma^\lambda = (\Gamma^\lambda(x_i))_{1 \leq i \leq N^\lambda}$ is discrete-time and vector-valued. Note that, this process only depends on the random selection of a relay for each transmitter which is done with probability $1/|Y^\lambda|$. For the process $\tilde{\Gamma}^\lambda$ the randomness comes from the uniform distribution of selection variables $U = (U_i)_{1 \leq i \leq N^\lambda}$.

For simplicity we assume the data transmission times to be ordered. Then, both $\Gamma^\lambda = \{\Gamma_{t_i}^\lambda\}$ and $\tilde{\Gamma}^\lambda = \{\tilde{\Gamma}_{t_i}^\lambda\}$ are time-discrete Markov chains in the finite state space $\{0, \lambda^{-1}\}^{N^\lambda}$. Moreover, at a given time t_i both, Γ^λ and $\tilde{\Gamma}^\lambda$ can only change in coordinate x_i. For the transition probability of Γ^λ we note that if at time t_i-, k relays are busy, then at time t_i, x_i becomes frustrated with probability $k/|Y^\lambda|$. Accordingly, with probability $1 - k/|Y^\lambda|$ at time t_i, x_i becomes satisfied and the process stays unchanged. As for $\tilde{\Gamma}^\lambda$ if at time t_i-, k relays are busy, then x_i becomes frustrated if $U_i \in [0, k/|Y^\lambda|]$. Hence the transition probabilities coincide. $\qquad\square$

2.2. *Existence and uniqueness of solutions.* For absolutely continuous measures $\mathcal{M}_{\mathsf{ac}} = \{\nu \in \mathcal{M} : \nu \ll \mu_\mathsf{T}\}$ the existence of solutions to (1.5) is non-trivial and will be dealt with in this section. More precisely, let $\mathcal{L} = \{f \in [0,1]^{[0,t_\mathsf{f}]} : f$ increasing and $f(0) = 0\}$, then for $\nu \in \mathcal{M}_{\mathsf{ac}}$ we define the integral operator $\mathcal{T}_\nu : \mathcal{L} \to \mathcal{L}$ where

$$\mathcal{T}_\nu : \beta \mapsto \Big(\int_0^t \nu(\mathrm{d}s, [\beta_s, 1], W) \Big)_{t \in [0, t_\mathsf{f}]}. \tag{2.1}$$

As will be shown below, $\mathcal{T}_\nu$ is continuous so that existence of solutions can be established via the Schauder-Tychonoff fixed point theorem. The uniqueness is a consequence of a monotonicity property of $\mathcal{T}_\nu$.

Proposition 2.2. *For all* $\nu \in \mathcal{M}_{\mathsf{ac}}$, *there exists exactly one* $\beta \in \mathcal{L}$ *such that* $\mathcal{T}_\nu(\beta) = \beta$.

Note that for the driving measure $\nu = L_\lambda$, existence and uniqueness of solutions for (1.5) is trivial. In both cases, where $\nu = L_\lambda$ or $\nu \in \mathcal{M}_{\mathsf{ac}}$, using (1.4), we obtain a measure-valued solution which we denote by $\beta(\nu)$. This solution is a step-function with step-height λ^{-1} if $\nu = L_\lambda$ and continuous if $\nu \in \mathcal{M}_{\mathsf{ac}}$.

2.3. *Approximation scheme for the solution.* In this subsection, we consider time-discretized two-step Picard approximations which will turn out to be convergent in the supremum norm. Additionally, the resulting trajectories exhibit good continuity

properties w.r.t. driving measures. Let us start by providing precise definitions of the approximations.

Let $\delta > 0$, $\nu \in \mathcal{M}$ and consider the discretization of $[0, t_{\mathsf{f}}]$ into t_{f}/δ disjoint segments of length δ, where we assume t_{f}/δ to be an integer. To define the approximation, we proceed recursively and start by putting $\beta_0^\delta(\nu)(\mathrm{d}x) = 0$. Once $\beta_{(n-1)\delta}^\delta(\nu)(\mathrm{d}x)$ is available, we define the locally constant function

$$\beta_t^{\uparrow,\delta}(\nu) = \beta_{(n-1)\delta}^\delta(\nu)(W) + \nu\big(((n-1)\delta, n\delta] \times [0,1] \times W\big)$$

for $t \in ((n-1)\delta, n\delta]$. Then, for $t \in ((n-1)\delta, n\delta]$ we put

$$\beta_t^\delta(\nu)(\mathrm{d}x) = \beta_{(n-1)\delta}^\delta(\nu)(\mathrm{d}x) + \nu\big(((n-1)\delta, t], [\beta_{n\delta}^{\uparrow,\delta}(\nu), 1], \mathrm{d}x\big).$$

We can think of $\beta^{\uparrow,\delta}(\nu)$ as a one-step Picard iteration of the zero function. Similarly, $\beta^\delta(\nu)$ corresponds to the two-step Picard iteration. The approximating process of frustrated transmitters is now defined as

$$\gamma_t^\delta(\nu)(\mathrm{d}x) = \nu([0,t], [0,1], \mathrm{d}x) - \beta_t^\delta(\nu)(\mathrm{d}x).$$

Next we show that $\gamma^\delta(\nu)$ approximates $\gamma(\nu)$ sufficiently well to transfer the LDP from $\gamma^\delta(\nu)$ to $\gamma(\nu)$. More precisely, we want to apply the exponential approximation technique of Eichelsbacher and Schmock (1998, Theorem 1.13) and therefore have to show three conditions. First, we show an exponential approximation property on the Banach space of trajectories of finite signed measures equipped with the supremum norm $\|\cdot\|$ where the supremum is taken over all times and measurable sets. The statement of this auxiliary result makes use of the empirical measures L_λ introduced below (1.3).

Proposition 2.3. $\gamma^\delta(L_\lambda)$ *is an* $\|\cdot\|$*-exponentially good approximation of* $\gamma(L_\lambda)$.

Second, we show the uniform approximation property on measures with bounded entropy.

Proposition 2.4. *For all* $\alpha \geq 0$, $\limsup_{\delta \downarrow 0} \sup_{\nu:\, h(\nu|\mu_{\mathsf{T}}) \leq \alpha} \|\gamma^\delta(\nu) - \gamma(\nu)\| = 0$.

Third, we show continuity of the approximation w.r.t. the driving measure in the τ-topology, i.e., the topology generated by evaluations on bounded measurable functions.

Proposition 2.5. *Let* $t \in [0, t_{\mathsf{f}}]$, $\delta > 0$ *and* $A \subset W$ *measurable, then at any* $\nu \in \mathcal{M}_{\mathsf{ac}}$, *the evaluations* $\mathcal{M} \to [0,\infty)$ *given by* $\nu \mapsto \gamma_t^\delta(\nu)(A)$ *are continuous w.r.t. the* τ*-topology.*

In particular, Propositions 2.4 and 2.5 imply that for all $\alpha > 0$, on $\nu:\, h(\nu|\mu_{\mathsf{T}}) \leq \alpha$, the map $\nu \mapsto \gamma(\nu)$ is continuous in the τ-topology.

Feeding the exponential approximation machinery with the Sanov-type LDP, we obtain the following multivariate LDP.

Proposition 2.6. *Let* $0 \leq t_1 \leq \cdots \leq t_k \leq t_{\mathsf{f}}$, *then the family of random measures* $\{\Gamma_{t_i}^\lambda(\mathrm{d}x)\}_i$ *satisfies the LDP in the* τ*-topology with good rate function* $I((\gamma_{t_i})_i) = \inf_{\nu \in \mathcal{M}:\, (\gamma_{t_i}(\nu))_i = (\gamma_{t_i})_i} h(\nu|\mu_{\mathsf{T}})$.

2.4. *Topological lifting and proof of Theorem 1.1.* In this section we first arrive at the continuous-path LDP using the Dawson-Gärtner Theorem (Dembo and Zeitouni, 1998, Theorem 4.6.1). We work in the product topology in $\mathcal{M}(W)^{[0,t_{\mathsf{f}}]}$, i.e., the coarsest topology such that the evaluations $\gamma \mapsto \gamma_t$ are continuous in the τ-topology. The topological lifting is then done using exponential tightness.

Proposition 2.7. *The family of measure-valued processes Γ^λ satisfies the LDP in the product topology. The good rate function is given by $I(\gamma) = \inf\limits_{\nu \in \mathcal{M}:\, \gamma(\nu) = \gamma} h(\nu|\mu_{\mathsf{T}})$.*

Proof of Theorem 1.1: First note that Proposition 2.7, by contraction, implies the same LDP where the τ-topology is replaced by the weak topology. Then, we may apply Dembo and Zeitouni (1998, Corollary 4.2.6) to reduce the proof of Theorem 1.1 to the exponential tightness of Γ^λ in the Skorohod topology. For this we use a criterion from Feng and Kurtz (2006, Theorem 4.1) and verify its conditions:

(1) Γ_t^λ is exponentially tight in $\mathcal{M}(W)$ for all $t \in [0, t_{\mathsf{f}}]$ and
(2) $\limsup_{\delta \downarrow 0} \limsup_{\lambda \uparrow \infty} \lambda^{-1} \log \mathbb{P}(w_\delta'(\Gamma^\lambda) > \varepsilon) = -\infty$.

Here the modulus of continuity w' is defined as

$$w_\delta'(\Gamma^\lambda) = \inf_{0 = t_0 < \cdots < t_k = t_{\mathsf{f}}:\, \min_{1 \leq i \leq k} |t_{i-1} - t_i| > \delta} \max_{1 \leq i \leq k} \sup_{s,t \in [t_{i-1}, t_i)} d_{\mathrm{w}}(\Gamma_s^\lambda, \Gamma_t^\lambda)$$

and with F^ε denoting the ε-halo of a closed set $F \subset W$,

$$d_{\mathrm{w}}(\nu, \nu') = \inf\{\varepsilon > 0 : \nu(F) \leq \nu'(F^\varepsilon) + \varepsilon,\ \nu'(F) \leq \nu(F^\varepsilon) + \varepsilon \text{ for all closed } F \subset W\}$$

is the Prokhorov metric on $\mathcal{M}(W)$ which makes $(\mathcal{M}(W), d_{\mathrm{w}})$ a Polish space, see Daley and Vere-Jones (1988, Proposition A2.5.III.).

As for (1) note that $K_\alpha = \{\nu \in \mathcal{M}(W) : \nu(W) \leq \alpha\}$ is compact in the weak topology for any $\alpha > 0$. Using the Poisson concentration inequality (Boucheron et al., 2013, Chapter 2.2),

$$\mathbb{P}(\Gamma_t^\lambda \in K_\alpha^c) = \mathbb{P}(\Gamma_t^\lambda(W) > \alpha) \leq \mathbb{P}(L_\lambda(V) > \alpha) \leq \exp(-\lambda h(\alpha|\mu_{\mathsf{T}}(V))),$$

where $h(x|y) = x \log(x/y) - x + y$. This shows the exponential tightness at every t.

For the exponential bound on the modulus of continuity (2) first note that $d_{\mathrm{w}}(\Gamma_s^\lambda, \Gamma_t^\lambda) \leq \sup_{1 \leq i \leq t_{\mathsf{f}}/\delta} L_\lambda([(i-1)\delta, i\delta] \times [0,1] \times W) = \rho$. Indeed, for all closed $F \subset W$ and $i\delta \leq s \leq t \leq (i+1)\delta$, it suffices to show that

$$\Gamma_t^\lambda(F) \leq \Gamma_s^\lambda(F^\rho) + \rho \tag{2.2}$$

since the other inequality $\Gamma_s^\lambda(F) \leq \Gamma_t^\lambda(F^\rho) + \rho$ is trivially satisfied for all $\rho > 0$. We can rewrite (2.2) equivalently as,

$$L_\lambda([s,t] \times [0,1] \times F)$$
$$\leq [B_t^\lambda(F) - B_s^\lambda(F)] + [L_\lambda([0,s] \times [0,1] \times F^\rho \setminus F) - B_s^\lambda(F^\rho \setminus F)] + \rho$$

where the first two summands on the r.h.s. are nonnegative. By our definition of ρ, we arrive at the desired bound. Consequently, using the Poisson concentration

inequality again,

$$\limsup_{\lambda\uparrow\infty} \lambda^{-1} \log \mathbb{P}(w'_\delta(\Gamma^\lambda) > \varepsilon)$$

$$\leq \limsup_{\lambda\uparrow\infty} \lambda^{-1} \log \mathbb{P}(\sup_{1\leq i\leq t_f/\delta} L_\lambda([(i-1)\delta, i\delta] \times [0,1] \times W) > \varepsilon)$$

$$\leq \limsup_{\lambda\uparrow\infty} \lambda^{-1} \log \sum_{1\leq i\leq t_f/\delta} \mathbb{P}(L_\lambda([(i-1)\delta, i\delta] \times [0,1] \times W) > \varepsilon)$$

$$\leq \max_{1\leq i\leq t_f/\delta} -h(\varepsilon|\mu_\mathsf{T}^\mathsf{s}(W)\mu_\mathsf{T}^\mathsf{t}([(i-1)\delta, i\delta]))$$

$$= -h(\varepsilon|\mu_\mathsf{T}^\mathsf{s}(W) \max_{1\leq i\leq t_f/\delta} \mu_\mathsf{T}^\mathsf{t}([(i-1)\delta, i\delta])).$$

Since $\mu_\mathsf{T}^\mathsf{t}$ is assumed to be absolutely-continuous w.r.t. the Lebesgue measure, this tends to minus infinity as δ tends to zero, as required. $\qquad\square$

3. Proofs of supporting results for Theorem 1.1

In this section, we provide proofs for the Propositions 2.2-2.7. To ease notation we will write in the following $\Delta_\delta(i) = ((i-1)\delta, i\delta]$. Let us start by stating three results that we will use multiple times in the sequel.

Lemma 3.1. *Let $\mathcal{B}(W) = \{A \subset W : A \text{ is Borel measurable}\}$ then the following holds.*

(1) *Let $\nu \in \mathcal{M}_\mathsf{ac}$, then*

$$\lim_{\varepsilon\downarrow 0} \sup_{A\subset\mathcal{B}(W):\, \mu_\mathsf{T}(A)<\varepsilon} \nu(A) = 0.$$

(2) *Let $\alpha > 0$, then*

$$\lim_{\varepsilon\downarrow 0} \sup_{\substack{A\subset\mathcal{B}(W):\, \mu_\mathsf{T}(A)<\varepsilon \\ \nu\in\mathcal{M}:\, h(\nu|\mu_\mathsf{T})<\alpha}} \nu(A) = 0.$$

(3) *Let $\delta > 0$ then, for a random variable $N^{\varepsilon\lambda}$ which is Poisson distributed with parameter $\varepsilon\lambda$*

$$\lim_{\varepsilon\downarrow 0} \limsup_{\lambda\uparrow\infty} \lambda^{-1} \log \mathbb{P}(N^{\varepsilon\lambda} > \lambda\delta) = -\infty.$$

Proof: Part 1 rephrases the definition of absolute continuity. Part 2 can be shown using Jensen's inequality. Part 3 is a consequence of the Poisson concentration inequality Boucheron et al. (2013, Chapter 2.2). $\qquad\square$

3.1. *Existence and uniqueness of solutions.* Let us start by establishing continuity of the integral operator.

Lemma 3.2. *Let $\nu \in \mathcal{M}_\mathsf{ac}$, then the map $\mathcal{T}_\nu : \mathcal{L} \to \mathcal{L}$ is continuous in the product topology.*

Proof: We need to show that the map $\beta \mapsto \mathcal{T}_\nu(\beta)_t$ is continuous for every $t \in [0, t_f]$. Observe that for any $\beta' \in \mathcal{L}$ and $\delta > 0$,

$$|\mathcal{T}_\nu(\beta')_t - \mathcal{T}_\nu(\beta)_t| \leq \sum_{i=1}^{t/\delta} \nu\big(\Delta_\delta(i) \times [\beta'_{(i-1)\delta} \wedge \beta_{(i-1)\delta}, \beta'_{i\delta} \vee \beta_{i\delta}] \times W\big).$$

For $\nu = \mu_{\mathsf{T}}$ the r.h.s. is bounded from above by

$$\mu_{\mathsf{T}}^{\mathsf{s}}(W) \sup_{1 \le i \le n} \mu_{\mathsf{T}}^{\mathsf{t}}(\Delta_\delta(i)) \sum_{i=1}^{n} (\beta'_{i\delta} \vee \beta_{i\delta} - \beta'_{(i-1)\delta} \wedge \beta_{(i-1)\delta}) \tag{3.1}$$

and

$$\sum_{i=1}^{t/\delta} (\beta'_{i\delta} \vee \beta_{i\delta} - \beta'_{(i-1)\delta} \wedge \beta_{(i-1)\delta}) \le 1 + 2 \sum_{i=1}^{t/\delta} |\beta'_{i\delta} - \beta_{i\delta}|.$$

Now, using Lemma 3.1 part 1, since $\mu_{\mathsf{T}}^{\mathsf{t}}$ is absolutely continuous w.r.t. the Lebesgue measure on $[0, t_{\mathsf{f}}]$, for any $\varepsilon > 0$, there exists $\delta > 0$ such that $\sup_{1 \le i \le t/\delta} \mu_{\mathsf{T}}^{\mathsf{t}}(\Delta_\delta(i)) < \varepsilon$. Secondly, for any fixed δ, by product-convergence, there exists a neighborhood of β such that if β' is in that neighborhood, $\sum_{i=1}^{t/\delta} |\beta'_{i\delta} - \beta_{i\delta}| \le 1/2$. In particular, (3.1) is bounded from above by $2\varepsilon\mu_{\mathsf{T}}^{\mathsf{s}}(W)$ and can be made arbitrarily small. Since $\nu \ll \mu_{\mathsf{T}}$ by assumption, using again Lemma 3.1 part 1, this transfers to ν and the proof is finished. $\qquad\square$

Using the above continuity, now existence and uniqueness follow from the Schauder-Tychonoff fixed-point theorem and monotonicity.

Proof of Proposition 2.2: Let us start by showing existence. Note that the Schauder-Tychonoff fixed-point theorem, see Granas and Dugundji (2003, Theorem II.7.1.10), implies existence if $\mathcal{T}_\nu : \mathcal{L} \to \mathcal{L}$ is continuous and $\mathcal{L}$ is a compact, convex subset of a locally convex linear topological space. For this first note that $\mathbb{R}^{[0,t_{\mathsf{f}}]}$ equipped with the product topology is a locally convex topological vector space. Further, note that $\mathcal{L}$ is closed inside the compact subset $[0,1]^{[0,t_{\mathsf{f}}]}$ and thereby compact. Since a convex combination of increasing functions is also increasing, $\mathcal{L}$ is also convex. By Lemma 3.2, the mapping $\beta \mapsto \mathcal{T}_\nu(\beta)$ is continuous, which implies existence.

As for the uniqueness, we proceed by contradiction, assuming that there exist two solutions $\beta, \beta' \in \mathcal{L}$ of (1.5) and a point in time $t_1 \in [0, t_{\mathsf{f}}]$ satisfying $\beta_{t_1} > \beta'_{t_1}$. Then, we let $t_0 \in [0, t_1)$ denote the last point before t_1, where $\beta_{t_0} = \beta'_{t_0}$. In particular,

$$\beta_{t_1} = \beta_{t_0} + \int_{(t_0,t_1]} \nu(\mathrm{d}s, [\beta_s, 1]) \le \beta'_{t_0} + \int_{(t_0,t_1]} \nu(\mathrm{d}s, [\beta'_s, 1]) = \beta'_{t_1},$$

which gives the desired contradiction. $\qquad\square$

3.2. *Exponential approximation property of the approximation scheme.* Let us first derive some results on dominance and closeness of the approximating trajectories w.r.t. the original process. We write Δ for the symmetric difference between sets.

Lemma 3.3. *Let $\delta > 0$ and $\nu \in \mathcal{M}_{\mathsf{ac}}$ or $\nu = L_\lambda$. Then, $\beta^\delta(\nu)(W) \le \beta(\nu)(W)$.*

Proof: We will abbreviate $\beta(\nu)(W) = \beta(\nu)$ and analogously for β^δ. It suffices to show that $\beta_t^\delta(\nu) \le \beta_t(\nu)$ holds for all $t \in \Delta_\delta(k)$ and $k \in \{0, \dots, t_{\mathsf{f}}/\delta\}$. We proceed by induction over k, the case $k = 0$ being trivial. Suppose that $\beta_{k\delta}^\delta(\nu) \le \beta_{k\delta}(\nu)$. In order to derive a contradiction, we assume that $\beta_t(\nu) < \beta_t^\delta(\nu)$ for some $t \in \Delta_\delta(k+1)$.

If $\nu \in \mathcal{M}_{\mathsf{ac}}$, there exists a largest time $t_1 \in [k\delta, t)$ such that $\beta_{t_1}(\nu) = \beta_{t_1}^\delta(\nu)$. In particular, for every $s \in (t_1, t)$

$$\beta_s^{\uparrow,\delta}(\nu) \ge \beta_s^\delta(\nu) \ge \beta_s(\nu)$$

and as required

$$\beta_t^\delta(\nu) = \beta_{t_1}^\delta(\nu) + \int_{(t_1,t]} \nu(\mathrm{d}s, [\beta_s^{\uparrow,\delta}(\nu), 1], W) \le \beta_{t_1}(\nu) + \int_{(t_1,t]} \nu(\mathrm{d}s, [\beta_s(\nu), 1], W)$$
$$= \beta_t(\nu).$$

If $\nu = L_\lambda$, there exists a largest time $t_1 \in [k\delta, t)$ such that $\beta_{t_1}^\delta(\nu) = \beta_{t_1}(\nu)$ and $\beta_{t_2}^\delta(\nu) > \beta_{t_2}(\nu)$, where t_2 is the next transmission time after t_1 in L_λ. In particular,

$$\beta_{t_2}^\delta(\nu) = \beta_{t_1}^\delta(\nu) + \nu(\{t_2\}, [\beta_{t_2}^{\uparrow,\delta}(\nu), 1], W)$$
$$\le \beta_{t_1}(\nu) + \nu(\{t_2\}, [\beta_{t_2}(\nu), 1], W) \le \beta_{t_2}(\nu),$$

as required. $\square$

The next lemma asserts an approximation property for scalar trajectories, uniform in time.

Lemma 3.4. *Assume $\nu \in \mathcal{M}_{\mathsf{ac}}$ or $\nu = L_\lambda$, then, for all $\delta > 0$,*

$$\|\beta^\delta(\nu)(W) - \beta(\nu)(W)\| \le 2 \sup_{1 \le i \le t_\mathsf{f}/\delta} \nu(\Delta_\delta(i) \times [0,1] \times W).$$

Proof: Again, we abbreviate $\beta(\nu)(W) = \beta(\nu)$ and analogously for β^δ. Let $\varepsilon = \sup_{1 \le i \le t_\mathsf{f}/\delta} \nu(\Delta_\delta(i) \times [0,1] \times W)$, then by Lemma 3.3 it suffices to show that

$$\beta_t(\nu) \le \beta_t^\delta(\nu) + 2\varepsilon$$

holds for all $t \in [0, t_\mathsf{f}]$. As before, we proceed by induction on the interval $\Delta_\delta(k)$ containing t. If $\beta_{(k-1)\delta}^\delta(\nu) + \varepsilon \ge \beta_{(k-1)\delta}(\nu)$, then

$$\beta_t(\nu) \le \beta_{(k-1)\delta}(\nu) + \nu(((k-1)\delta, t] \times [0,1] \times W) \le (\beta_{(k-1)\delta}^\delta(\nu) + \varepsilon) + \varepsilon \le \beta_t^\delta(\nu) + 2\varepsilon.$$

Otherwise, if $\beta_{(k-1)\delta}^\delta(\nu) + \varepsilon \le \beta_{(k-1)\delta}(\nu)$, then

$$\beta_{k\delta}^{\uparrow,\delta}(\nu) = \beta_{(k-1)\delta}^\delta(\nu) + \nu(\Delta_\delta(k) \times [0,1] \times W) \le \beta_{(k-1)\delta}(\nu).$$

Hence,

$$\beta_t(\nu) = \beta_{(k-1)\delta}(\nu) + \int_{(k-1)\delta}^t \nu(\mathrm{d}s, [\beta_{s-}(\nu), 1], W)$$

$$\le \beta_{(k-1)\delta}(\nu) + \nu(((k-1)\delta, t] \times [\beta_{k\delta}^{\uparrow,\delta}(\nu), 1], W) = \beta_{(k-1)\delta}(\nu) + \beta_t^\delta(\nu) - \beta_{(k-1)\delta}^\delta(\nu),$$

so that the assertion follows from the induction hypothesis. $\square$

We denote in the sequel $I_t^\delta(\nu) = [\beta_t(\nu)(W) \wedge \beta_t^{\uparrow,\delta}(\nu), \beta_t(\nu)(W) \vee \beta_t^{\uparrow,\delta}(\nu)]$ and note the following approximation property for measure-valued trajectories, uniform in time and over measurable sets. For $\nu \in \mathcal{M}_{\mathsf{ac}}$ or $\nu = L_\lambda$ and for all $\delta > 0$,

$$\|\beta^\delta(\nu) - \beta(\nu)\| \le \int_0^t \nu(\mathrm{d}s, [\beta_{s-}(\nu), 1]\Delta[\beta_{s-}^{\uparrow,\delta}(\nu), 1], W) \le \int_0^{t_\mathsf{f}} \nu(\mathrm{d}t, I_{t-}^\delta(\nu), W).$$
$$(3.2)$$

We show that in the setting of empirical measures, the inequality (3.2) gives rise to a strong probabilistic bound on the $\|\cdot\|$-distance. First, we need an auxiliary result on stochastic domination.

Lemma 3.5. *Let $\varepsilon > 0$ be arbitrary. Then, the random variable $\int_0^{t_f} X^\lambda(\mathrm{d}t, [B^\lambda_{t-}(W) - \varepsilon, B^\lambda_{t-}(W) + \varepsilon], W)$ is stochastically dominated by a Poisson random variable with parameter $2\varepsilon\mu^s_\mathsf{T}(W)$.*

Proof: The proof is based on the previsiblity of the time integral. More precisely, let $\rho \geq 0$ and recall that conditioned on $|X^\lambda| = n$ the marks $(U_i)_{1 \leq i \leq n}$ are iid. Ordering X^λ by the time-index yields

$$\mathbb{P}\Big(\int_0^{t_f} X^\lambda(\mathrm{d}t, [B^\lambda_{t-}(W) - \varepsilon, B^\lambda_{t-}(W) + \varepsilon], W) > \rho\big||X^\lambda| = n\Big)$$

$$= \mathbb{E}\Big[\mathbb{P}\big(\mathbb{1}_{\{|B_{t_{n-1}}(W) - U_n| \leq \varepsilon\}} + \sum_{i=1}^{n-1} \mathbb{1}_{\{|B_{t_{i-1}}(W) - U_i| \leq \varepsilon\}} > \rho\big|(X^\lambda_i)_{1 \leq i \leq n-1}\big)\big||X^\lambda| = n\Big]$$

$$\leq \mathbb{P}\Big(1_{[0,2\varepsilon]}(U_n) + \sum_{i=1}^{n-1} \mathbb{1}\{|B_{t_{i-1}}(W) - U_i| \leq \varepsilon\} > \rho\big||X^\lambda| = n\Big)$$

$$\leq \mathbb{P}\Big(\sum_{i=1}^{n} 1_{[0,2\varepsilon]}(U_i) > \rho\big||X^\lambda| = n\Big).$$

In particular,

$$\mathbb{P}\Big(\int_0^{t_f} X^\lambda(\mathrm{d}t, [B^\lambda_{t-}(W) - \varepsilon, B^\lambda_{t-}(W) + \varepsilon], W) > \rho\Big) \leq \mathbb{P}\Big(\sum_{i=1}^{|X^\lambda|} 1_{[0,2\varepsilon]}(U_i) > \rho\Big)$$

and by independent thinning, $\sum_{i=1}^{|X^\lambda|} 1_{[0,2\varepsilon]}(U_i)$ is a Poisson random variable with the desired parameter. $\qquad\square$

Note that, by the definition of γ and γ^δ, we have

$$\gamma^\delta(\nu) - \gamma(\nu) = \beta^\delta(\nu) - \beta(\nu).$$

Therefore in the proofs of Proposition 2.3, 2.4 and 2.5, γ and γ^δ can be replaced by β and β^δ.

Proof of Proposition 2.3: Let $\varepsilon > 0$ be arbitrary. By the definition of exponential good approximations (Eichelsbacher and Schmock, 1998, Definition 1.2), we need to check that

$$\limsup_{\delta\downarrow0} \limsup_{\lambda\uparrow0} \lambda^{-1} \log \mathbb{P}(\|B^\lambda - \beta^\delta(L_\lambda)\| > \varepsilon) = -\infty.$$

Using the bound (3.2), we have for all $\varepsilon' > 0$ the uniform estimate

$$\mathbb{P}(\|B^\lambda - \beta^\delta(L_\lambda)\| > \varepsilon) \leq \mathbb{P}\Big(\int_0^{t_f} L_\lambda(\mathrm{d}t, I^\delta_{t-}(L_\lambda), W) > \varepsilon\Big)$$

$$\leq \mathbb{P}\Big(\int_0^{t_f} L_\lambda(\mathrm{d}t, [B^\lambda_{t-}(W) - \varepsilon', B^\lambda_{t-}(W) + \varepsilon'], W) > \varepsilon\Big) + \mathbb{P}(\sup_{t\in[0,t_f]} |I^\delta_t(L_\lambda)| > \varepsilon').$$
$$\tag{3.3}$$

For the first summand on the r.h.s. of (3.3) we can use Lemmas 3.5 and 3.1 part 3. For the second summand on the r.h.s. of (3.3), note that

$$\mathbb{P}(\sup_{t\in[0,t_f]} |I^\delta_t(L_\lambda)| > \varepsilon')$$

$$\leq \mathbb{P}(\|\beta^\delta(L_\lambda)(W) - \beta^{\uparrow,\delta}(L_\lambda)\| > \varepsilon'/2) + \mathbb{P}(\|B^\lambda(W) - \beta^\delta(L_\lambda)(W)\| > \varepsilon'/2).$$

By definition, respectively by Lemma 3.4, we have

$$\mathbb{P}(\|\beta^\delta(L_\lambda)(W) - \beta^{\uparrow,\delta}(L_\lambda)\| > \varepsilon'/2) \leq \mathbb{P}(\sup_{1 \leq i \leq t_\mathrm{f}/\delta} L_\lambda(\Delta_\delta(i) \times [0,1] \times W) > \varepsilon'/2)$$

$$\mathbb{P}(\|B^\lambda(W) - \beta^\delta(L_\lambda)(W)\| > \varepsilon'/2) \leq \mathbb{P}(\sup_{1 \leq i \leq t_\mathrm{f}/\delta} L_\lambda(\Delta_\delta(i) \times [0,1] \times W) > \varepsilon'/4).$$

Using again Lemma 3.1 part 3, the proof is finished. $\qquad\square$

Proof of Proposition 2.4: By the bound (3.2) we have the estimate

$$\|\beta(\nu) - \beta^\delta(\nu)\| \leq \int_0^{t_\mathrm{f}} \nu(\mathrm{d}t, I_t^\delta(\nu), W).$$

Moreover, by Lemma 3.4 and the definition of $\beta^{\uparrow,\delta}$,

$$\sup_{t \in [0,t_\mathrm{f}]} |I_t^\delta(\nu)| \leq 2 \sup_{1 \leq i \leq t_\mathrm{f}/\delta} \nu(\Delta_\delta(i) \times [0,1] \times W).$$

It follows by Lemma 3.1 parts 1 and 2 that

$$\limsup_{\delta \downarrow 0} \sup_{\substack{t \in [0,t_\mathrm{f}] \\ \nu:\, h(\nu|\mu_\mathsf{T}) \leq \alpha}} |I_t^\delta(\nu)| = 0.$$

Consequently, for all $\varepsilon > 0$ and sufficiently small $\delta > 0$,

$$\limsup_{\delta \downarrow 0} \sup_{\nu:\, h(\nu|\mu_\mathsf{T}) \leq \alpha} \int_0^{t_\mathrm{f}} \nu(\mathrm{d}t, I_t^\delta(\nu), W)$$

$$\leq \limsup_{\varepsilon \downarrow 0} \sup_{\nu:\, h(\nu|\mu_\mathsf{T}) \leq \alpha} \int_0^{t_\mathrm{f}} \nu(\mathrm{d}t, [\beta_t(\nu)(W) - \varepsilon, \beta_t(\nu)(W) + \varepsilon], W).$$

Another application of Lemma 3.1 parts 1 and 2 gives the result. $\qquad\square$

Proof of Proposition 2.5: Assume $\nu' \in \mathcal{M}$, $\nu \in \mathcal{M}_\mathsf{ac}$ and consider $|\beta_t^\delta(\nu')(A) - \beta_t^\delta(\nu)(A)|$ for some $t \in [0,t_\mathrm{f}]$ and measurable $A \subset W$. Then for $n \in \{1, \ldots, t_\mathrm{f}/\delta\}$ such that $t \in \Delta_\delta(n)$ we have the upper bound

$$|\beta_t^\delta(\nu')(A) - \beta_t^\delta(\nu)(A)| \leq |\beta_{(n-1)\delta}^\delta(\nu')(A) - \beta_{(n-1)\delta}^\delta(\nu)(A)|$$
$$+ |\nu'(((n-1)\delta, t] \times [\beta_{n\delta}^{\uparrow,\delta}(\nu'), 1] \times A) - \nu(((n-1)\delta, t] \times [\beta_{n\delta}^{\uparrow,\delta}(\nu), 1] \times A)|.$$

Using the estimate

$$|\beta_{i\delta}^\delta(\nu')(A) - \beta_{i\delta}^\delta(\nu)(A)| \leq |\beta_{(i-1)\delta}^\delta(\nu')(A) - \beta_{(i-1)\delta}^\delta(\nu)(A)|$$
$$+ |\nu'(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow,\delta}(\nu'), 1] \times A) - \nu(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow,\delta}(\nu), 1] \times A)|,$$

we can further bound $|\beta_{(n-1)\delta}^\delta(\nu')(A) - \beta_{(n-1)\delta}^\delta(\nu)(A)|$ from above by

$$\sum_{i=1}^{t_\mathrm{f}/\delta} |\nu'(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow,\delta}(\nu'), 1] \times A) - \nu(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow,\delta}(\nu), 1] \times A)|.$$

We will suppress the spatial component A in our notation for the rest of the proof. Since the sum is finite, it suffices to consider any $1 \leq i \leq t_\mathrm{f}/\delta$ and note that the case where $((i-1)\delta, i\delta]$ is replaced by $((n-1)\delta, t]$ works equivalently. We can further estimate,

$$|\nu'(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow,\delta}(\nu'), 1]) - \nu(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow,\delta}(\nu), 1])|$$
$$\leq \nu'(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow,\delta}(\nu'), 1] \Delta [\beta_{i\delta}^{\uparrow,\delta}(\nu), 1]) + |(\nu' - \nu)(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow,\delta}(\nu), 1])|.$$

Now, for ν' sufficiently close to ν in the τ-topology, the second summand can be made arbitrarily small. Let $\varepsilon > 0$, then it suffices to show that for all ν' in a neighborhood of ν we have

$$\nu'(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow;\delta}(\nu'), 1]\Delta[\beta_{i\delta}^{\uparrow;\delta}(\nu), 1]) < \varepsilon.$$

For this, note that for all ε', there exists a neighborhood of ν such that for all ν' in that neighborhood

$$|\beta_{i\delta}^{\uparrow;\delta}(\nu') - \beta_{i\delta}^{\uparrow;\delta}(\nu)| \leq \sum_{j=1}^{t_f/\delta} |(\nu' - \nu)(\Delta_\delta(j) \times [0, 1])| < \varepsilon'.$$

For such ν', we thus have

$$\nu'(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow;\delta}(\nu'), 1]\Delta[\beta_{i\delta}^{\uparrow;\delta}(\nu), 1]) \leq \nu'(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow;\delta}(\nu) - \varepsilon', \beta_{i\delta}^{\uparrow;\delta}(\nu) + \varepsilon']).$$
(3.4)

Applying the definition of τ-convergence for the third time, for ν' sufficiently close to ν, the r.h.s. of (3.4) is close to

$$\nu(\Delta_\delta(i) \times [\beta_{i\delta}^{\uparrow;\delta}(\nu) - \varepsilon', \beta_{i\delta}^{\uparrow;\delta}(\nu) + \varepsilon'])$$

up to an arbitrarily small error. Finally, applying Lemma 3.1 part 1, the proof is finished. □

3.3. *Sanov's theorem and proof of Proposition 2.6.* For a sequence of iid random variables, Sanov's theorem in the τ-topology is one of the cornerstones of large deviations theory. Clearly, this result should remain valid when passing from the iid to the Poisson setting. However, as it is not easy to find a reference, we provide a detailed proof along the Gärtner-Ellis type argumentation presented in Dembo and Zeitouni (1998, Section 6.2). In our presentation, we focus on the steps where there is a substantial difference between the Poisson and the iid case. For the convenience of the reader, we adapt the notation from Dembo and Zeitouni (1998, Section 6.2) where possible.

Proposition 3.6. *The random measures L_λ satisfy the LDP in the τ-topology with good rate function given by*

$$I(\nu) = h(\nu|\mu_\mathsf{T}).$$

Moreover, the levelsets of I are sequentially compact in the τ-topology.

Proof: The empirical measure L_λ is a random variable with values in the space $\mathcal{X} = B(V)'$, the algebraic dual of the space of all bounded linear functions on V. We consider $\mathcal{X}$ as a vector space endowed with the topology generated by the evaluations $\nu \mapsto \nu(\varphi)$, $\varphi \in B(V)$. With this topology, the topological dual $\mathcal{X}^*$ of $\mathcal{X}$ is isomorphic to $B(V)$. Since the Laplace functional of a Poisson point process is known in closed form, the limiting logarithmic moment generating function of L_λ can be computed explicitly and is given by

$$\Lambda(\varphi) = \int [\exp(\varphi(v)) - 1]\mu_\mathsf{T}(\mathrm{d}v), \qquad \varphi \in B(V).$$

Since for every $\varphi_1, \ldots, \varphi_n \in B(V)$ the function $(t_1, \ldots, t_n) \mapsto \Lambda(\sum_i t_i \varphi_i)$ is everywhere differentiable, Dembo and Zeitouni (1998, Corollary 4.6.11) implies that L_λ satisfies the LDP with good rate function given by the Legendre dual Λ^* of Λ.

It remains to show that $\Lambda^* = h(\cdot|\mu_{\mathsf{T}})$. By duality theory (Dembo and Zeitouni, 1998, Lemmas 4.5.8 and 6.2.16), it suffices to show that

$$\Lambda(\varphi) = \sup_{\nu \in \mathcal{X}} \{\nu(\varphi) - h(\nu|\mu_{\mathsf{T}})\}. \tag{3.5}$$

In order to show that $\Lambda(\varphi) \leq \sup_{\nu \in \mathcal{X}} \{\nu(\varphi) - h(\nu|\mu_{\mathsf{T}})\}$ let ν_φ be the measure with density e^φ w.r.t. μ_{T}. Then, a quick computation shows that $\Lambda(\varphi) = \nu_\varphi(\varphi) - h(\nu_\varphi|\mu_{\mathsf{T}})$. Conversely, the r.h.s. of (3.5) is equal to $\sup_{\nu \in \mathcal{X}^*} \{\Lambda(\varphi) - h(\nu|\nu_\varphi)\}$ and the non-negativity of the entropy concludes the identification. Sequential compactness of $h(\cdot|\mu_{\mathsf{T}})$ follows from Dembo and Zeitouni (1998, Lemma 6.2.16). □

Recall that Γ^λ is constructed from the solution of (1.4) with $r = r_\lambda$. We will sometimes make this dependence explicit by writing

$$\gamma(L_\lambda, r_\lambda) = \Gamma^\lambda.$$

The results collected so far would allow us to derive the LDP similar to the one in Proposition 2.6 where Γ^λ is replaced by $\gamma(L_\lambda, 1)$. In order to conclude we thus need a final result on the asymptotic contribution of r_λ. Let us start with the following dominance result, where we write

$$B^{\lambda,r}(\mathrm{d}x) = L_\lambda([0,t],[0,1],\mathrm{d}x) - \gamma(L_\lambda, r)(\mathrm{d}x).$$

Lemma 3.7. *If $s \leq r$, then*

$$\tfrac{s}{r}(B^{\lambda,r}(W) - \lambda^{-1}) \leq B^{\lambda,s}(W) \leq B^{\lambda,r}(W).$$

Proof: We prove both claims by induction on the arrival time. Hence, we assume that the desired inequalities are valid up to time t_{i-1}. Now, suppose that at time t_i we had

$$B_{t_i}^{\lambda,s}(W) > B_{t_i}^{\lambda,r}(W).$$

This is only possible if $B_{t_{i-1}}^{\lambda,s}(W) = B_{t_{i-1}}^{\lambda,r}(W)$ and for the i'th choice variable U_i, drawn at time t_i, we have $U_i \in [s^{-1}B_{t_{i-1}}^{\lambda,s}(W), r^{-1}B_{t_{i-1}}^{\lambda,r}(W)]$. But since $r \leq s$, this is impossible.

Similarly, assume that at time t_i we had

$$\tfrac{s}{r}(B_{t_i}^{\lambda,r}(W) - \lambda^{-1}) > B_{t_i}^{\lambda,s}(W).$$

This is only possible if for the i'th choice variable $U_i \in [r^{-1}B_{t_{i-1}}^{\lambda,r}(W), s^{-1}B_{t_{i-1}}^{\lambda,s}(W)]$, so that $\tfrac{s}{r}B_{t_{i-1}}^{\lambda,r}(W) \leq B_{t_{i-1}}^{\lambda,s}(W)$. But this implies that

$$\tfrac{s}{r}(B_{t_i}^{\lambda,r}(W) - \lambda^{-1}) = \tfrac{s}{r}B_{t_{i-1}}^{\lambda,r}(W) \leq B_{t_{i-1}}^{\lambda,s}(W) = B_{t_i}^{\lambda,s}(W),$$

yielding the desired contradiction. □

Proposition 3.8. *The families of measure-valued processes Γ^λ and $\gamma(L_\lambda, 1)$ are $\|\cdot\|$-exponentially equivalent.*

Proof: Let $\varepsilon > 0$ be arbitrary. By the definition of exponential equivalence (Dembo and Zeitouni, 1998, Definition 4.2.14) and the identity $\|\gamma(L_\lambda, r_\lambda) - \gamma(L_\lambda, 1)\| = \|B^{\lambda,r_\lambda} - B^{\lambda,1}\|$, we need to check that

$$\limsup_{\lambda \uparrow \infty} \lambda^{-1} \log \mathbb{P}(\|B^{\lambda,r_\lambda} - B^{\lambda,1}\| > \varepsilon) = -\infty.$$

Arguing similarly as in the proof of Proposition 2.3, for any $\rho > 0$, the following estimate holds,

$$\mathbb{P}(\|B^{\lambda,r_\lambda} - B^{\lambda,1}\| > \varepsilon) \leq \mathbb{P}\Big(\int_0^{t_f} L_\lambda(dt, [\tfrac{B^{\lambda,r_\lambda}_{t-}(W)}{r_\lambda}, 1]\Delta[B^{\lambda,1}_{t-}(W), 1], W) > \varepsilon\Big)$$

$$\leq \mathbb{P}\Big(\int_0^{t_f} L_\lambda(dt, [B^{\lambda,1}_{t-}(W) - \rho, B^{\lambda,1}_{t-}(W) + \rho], W) > \varepsilon\Big) \tag{3.6}$$

$$+ \mathbb{P}(\|B^{\lambda,1}(W) - \tfrac{B^{\lambda,r_\lambda}(W)}{r_\lambda}\| > \rho).$$

For the first summand on the r.h.s. of (3.6) we can use Lemmas 3.5 and 3.1 part 3. For the second summand on the r.h.s. of (3.6) we can further estimate

$$\mathbb{P}(\|B^{\lambda,1}(W) - \tfrac{B^{\lambda,r_\lambda}(W)}{r_\lambda}\| > \rho) \leq \mathbb{P}(\|B^{\lambda,1}(W) - B^{\lambda,r_\lambda}(W)\| > \rho/2)$$

$$+ \mathbb{P}(B^{\lambda,r_\lambda}_{t_f}(W) > \tfrac{\rho r_\lambda}{2|1-r_\lambda|}).$$

As for the second summand using Lemma 3.1 part 3 we arrive at the desired limit. For the first summand, Lemma 3.7, implies that

$$\mathbb{P}(\|B^{\lambda,r_\lambda}(W) - B^{\lambda,1}(W)\| > \rho/2) \leq \mathbb{P}(B^{\lambda,1\vee r_\lambda}_{t_f}(W) > \tfrac{\rho/2 - \lambda^{-1}}{|1-r_\lambda|}),$$

so that the desired limit follows from Lemma 3.1 part 3. $\square$

Proof of Proposition 2.6: We first use Proposition 3.8 to remove the λ-dependence in the relays. To conclude the proof, we then apply the τ-topology version of the exponential approximation machinery (Eichelsbacher and Schmock, 1998, Theorem 1.13). The conditions are satisfied according to the Propositions 2.3, 2.4, 2.5 and 3.6. $\square$

3.4. *Dawson-Gärtner for the temporal component.* In order to derive the LDP for the product topology in the time dimension, we use Proposition 2.6 and apply the Dawson-Gärtner Theorem (Dembo and Zeitouni, 1998, Theorem 4.6.1).

Proof of Proposition 2.7: Consider the time-discretizations $\mathbf{t} = \{(t_0, \ldots, t_n) : 0 = t_0 < t_1 < \cdots < t_n = t_f\}$ and the associated projections $p_\mathbf{t}(\gamma) = (\gamma_{t_i})_{t_i \in \mathbf{t}}$. The family of all such projections J has a partial ordering induced by inclusion in the family of discretizations $\mathbf{t}$. Using Proposition 2.6 and Dembo and Zeitouni (1998, Theorem 4.6.1), Γ^λ satisfies the LDP in the product topology with good rate function given by

$$\tilde{I}(\gamma) = \sup_{\mathbf{t} \in J} I_\mathbf{t}(p_\mathbf{t}(\gamma)) \quad \text{where} \quad I_\mathbf{t}((\gamma_{t_i})_{t_i \in \mathbf{t}}) = \inf_{\nu \in \mathcal{M} : (\gamma_{t_i}(\nu))_{t_i \in \mathbf{t}} = (\gamma_{t_i})_{t_i \in \mathbf{t}}} h(\nu|\mu_\mathsf{T}).$$

The proof is finished once we show $\tilde{I}(\gamma) = I(\gamma)$ with $I(\gamma) = \inf_{\nu \in \mathcal{M} : \gamma(\nu) = \gamma} h(\nu|\mu_\mathsf{T})$.

We first prove $I \geq \tilde{I}$. Let $\nu' \in \mathcal{M}$ be such that $\gamma(\nu') = \gamma$. Then, in particular $(\gamma_{t_i}(\nu'))_{t_i \in \mathbf{t}} = (\gamma_{t_i})_{t_i \in \mathbf{t}}$ for any time-discretization $\mathbf{t}$ and

$$\inf_{\nu \in \mathcal{M} : (\gamma_{t_i}(\nu))_{t_i \in \mathbf{t}} = (\gamma_{t_i})_{t_i \in \mathbf{t}}} h(\nu|\mu_\mathsf{T}) \leq h(\nu'|\mu_\mathsf{T}),$$

so that $\tilde{I} \leq I$.

For the other direction, $I \leq \tilde{I}$, first assume that γ is discontinuous in the sense that there exists a measurable set $A \subset W$ and some $t_d \in [0, t_f]$ such that there exists a sequence $t_n \to t_d$ with $\lim_{n\to\infty} \gamma_{t_n}(A) \neq \gamma_{t_d}(A)$. Then, we show that

$\tilde{I}(\gamma) = \infty$ which trivially implies the inequality. Indeed, consider the sequence of time-partitions $\mathbf{t}_n = \{0 < t_n < t_d < t_{t_f}\}$. Then, there exists a sequence $\nu_n \in \mathcal{M}$ such that

$$\tilde{I}(\gamma) \geq \limsup_{n\uparrow\infty} \inf_{\nu\in\mathcal{M}:\, (\gamma_{t_i}(\nu))_{t_i\in\mathbf{t}_n}=(\gamma_{t_i})_{t_i\in\mathbf{t}_n}} h(\nu|\mu_{\mathsf{T}})$$

$$\geq \limsup_{n\uparrow\infty} \inf_{\nu\in\mathcal{M}:\, (\gamma_{t_i}(\nu)(A))_{t_i\in\mathbf{t}_n}=(\gamma_{t_i}(A))_{t_i\in\mathbf{t}_n}} h(\nu|\mu_{\mathsf{T}}) \geq \limsup_{n\uparrow\infty} h(\nu_n|\mu_{\mathsf{T}}) - \varepsilon$$

and $\gamma_t(\nu_n)(A) = \gamma_t(A)$ for all $t \in \mathbf{t}_n$. Moreover, setting $A_n = (t_n, t_d] \times [0,1] \times A$ then $\lim_{n\uparrow\infty} \mu_{\mathsf{T}}(A_n) = 0$ since $\mu_{\mathsf{T}}^{\mathsf{t}}$ is absolutely-continuous w.r.t. the Lebesgue measure. On the other hand, by assumption, there exists an $\varepsilon > 0$ such that

$$\varepsilon < |\gamma_{t_n}(A) - \gamma_{t_d}(A)| = |\gamma_{t_n}(\nu_n)(A) - \gamma_{t_d}(\nu_n)(A)|$$

for sufficiently large n. Hence $\nu_n(A_n) > \varepsilon/2$, so that Lemma 3.1 part 2 yields $\tilde{I}(\gamma) = \infty$.

It remains to consider the setting, where γ is continuous and $\tilde{I}(\gamma) < \infty$. Let $\mathbf{t}_\delta$ denote a finite partition of $[0, t_f]$ with mesh size smaller than $\delta > 0$. Then again, there exists a sequence $\nu_\delta \in \mathcal{M}$ such that

$$\tilde{I}(\gamma) \geq \limsup_{\delta\downarrow 0} \inf_{\nu\in\mathcal{M}:\, (\gamma_{t_i}(\nu))_{t_i\in\mathbf{t}_\delta}=(\gamma_{t_i})_{t_i\in\mathbf{t}_\delta}} h(\nu|\mu_{\mathsf{T}}) \geq \limsup_{\delta\downarrow 0} h(\nu_\delta|\mu_{\mathsf{T}}) - \varepsilon$$

and $\gamma_t(\nu_\delta) = \gamma_t$ for all $t \in \mathbf{t}_\delta$. Since the levelsets of $h(\cdot|\mu_{\mathsf{T}})$ are sequentially compact in the τ-topology, there exists a τ-accumulation point ν_* for $(\nu_\delta)_\delta$ and by the time-continuity of γ and the continuity of $\nu \mapsto \gamma(\nu)$ along sequences of measures with uniformly bounded entropy, we have $\gamma_t(\nu_*) = \gamma_t$ for all $t \in [0, t_f]$. Moreover, by the lower semicontinuity of $h(\cdot|\mu_{\mathsf{T}})$ we have $\limsup_{\delta\downarrow 0} h(\nu_\delta|\mu_{\mathsf{T}}) \geq h(\nu_*|\mu_{\mathsf{T}}) \geq I(\gamma)$. This finishes the proof. $\qquad\square$

4. Proof of Theorem 1.2

For the proof of Theorem 1.2 we construct spatially approximating processes by replacing κ with a carefully chosen step function. As for the time approximation considered in the proof of Theorem 1.1, the strongly regularizing property of the differential equation allows us to verify that again the approximation is uniformly close and exponentially approximates the original process. This reveals a striking methodological similarity between the space approximations appearing in the proof of Theorem 1.2 and the time approximations considered in Section 2. To implement this program, we first need to overcome the technical obstacle that the step functions still depend on the empirical relay process. In particular, Theorem 1.1 cannot yet be applied. A preliminary step is therefore to replace the relay process by its limiting measure and show that the error made is exponentially small.

4.1. *Exponential equivalence w.r.t. the relay process.* Let $W^\delta = \{W_1, \ldots, W_k\}$ be a partition of W into cubes of side length δ. If W is not a cube itself, then the W_i are defined as the intersection of the smaller cubes with W. The idea is to partition the transmitter process into independent processes, each process confined to choose relays in a given spatial discretization. More precisely, recalling (1.7), let $Z^{\lambda,i}(\nu_{\mathsf{R}})$ denote the Poisson point process with intensity measure

$$\mu^i(\nu_{\mathsf{R}})(\mathrm{d}s, \mathrm{d}u, \mathrm{d}x) = \kappa_{\nu_{\mathsf{R}}}(W_i|x)\mu_{\mathsf{T}}(\mathrm{d}s, \mathrm{d}u, \mathrm{d}x)$$

and let $L^i_\lambda(\nu_\mathsf{R})$ be the associated empirical measure. In other words, $Z^{\lambda,i}(\nu_\mathsf{R})$ is the Poisson point process of transmitters choosing a relay in W_i. Now, consider an associated augmented empirical measure given by

$$\mathfrak{L}^\delta_\lambda(\nu_\mathsf{R}) = \sum_{i=1}^k \mathfrak{L}^i_\lambda(\nu_\mathsf{R})$$

where

$$\mathfrak{L}^i_\lambda(\nu_\mathsf{R}) = L^i_\lambda(\nu_\mathsf{R}) \otimes 1_{W_i} \frac{\nu_\mathsf{R}}{\nu_\mathsf{R}(W_i)}.$$

Note that $\mathfrak{L}^\delta_\lambda(\nu_\mathsf{R})(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, W_i) = L^i_\lambda(\nu_\mathsf{R})(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x)$, so that the total mass of transmitters pointing into W_i is preserved. However, within W_i this mass is now distributed according to ν_R conditioned on W_i. In particular, the kernel $y \mapsto \mathfrak{L}^\delta_\lambda(\nu_\mathsf{R})_y$ appearing in (1.9) is constant on W_i, where it is given by $\nu_\mathsf{R}(W_i)^{-1} L^i_\lambda(\nu_\mathsf{R})$. Thus,

$$\int_{W_i} \gamma(\mathfrak{L}^\delta_\lambda(\nu_\mathsf{R})_y)\nu_\mathsf{R}(\mathrm{d}y) = \int_{W_i} \gamma(\nu_\mathsf{R}(W_i)^{-1} L^i_\lambda(\nu_\mathsf{R}))\nu_\mathsf{R}(\mathrm{d}y) = \gamma(L^i_\lambda(\nu_\mathsf{R}), \nu_\mathsf{R}(W_i)),$$

(4.1)

where we recall the more detailed notation $\gamma(\cdot, \cdot)$ from the paragraph preceding Lemma 3.7. In the proof of Theorem 1.2 this identification is an essential ingredient to establish a connection to the setting of Theorem 1.1.

As a first step, we show that it is possible to switch between $\nu_\mathsf{R} = l_\lambda$ and $\nu_\mathsf{R} = \mu_\mathsf{R}$ without changing substantially the approximating process of frustrated transmitters.

Proposition 4.1. *The family of measure-valued processes $\gamma(\mathfrak{L}^\delta_\lambda(l_\lambda), l_\lambda)$ is $\|\cdot\|$-exponentially equivalent to $\gamma(\mathfrak{L}^\delta_\lambda(\mu_\mathsf{R}), \mu_\mathsf{R})$.*

Next we show that $\gamma(\mathfrak{L}^\delta_\lambda(l_\lambda), l_\lambda)$ is an exponentially good approximation to Γ^λ.

Proposition 4.2. *The family of measure-valued processes $\gamma(\mathfrak{L}^\delta_\lambda(l_\lambda), l_\lambda)$ is an $\|\cdot\|$-exponentially good approximations of Γ^λ.*

Combining this with a uniform bound on the spatial discretization, we arrive at the following multivariate LDP.

Proposition 4.3. *Let $0 \le t_1 \le \cdots \le t_k \le t_\mathsf{f}$, then the family of random measures $\{\Gamma^\lambda_{t_i}(\mathrm{d}x)\}_i$ satisfies the LDP in the τ-topology with good rate function*

$$I((\gamma_{t_i})_i) = \inf_{\mathfrak{n} \in \mathcal{M}' : (\gamma_{t_i}(\mathfrak{n}, \mu_\mathsf{R}))_i = (\gamma_{t_i})_i} h(\mathfrak{n}|\mu(\mu_\mathsf{R})).$$

Proof of Theorem 1.2: Using a slight modification of the proof of Proposition 2.7, Proposition 4.2 can be lifted to the same LDP w.r.t. continuous times in the product topology. In order to finally establish the LDP in the Skorohod topology, the exponential tightness arguments presented in the proof of Theorem 1.1 should be applied verbatim.   □

5. Proof of supporting results for Theorem 1.2

In the previous section, we announced our plan to prove Theorem 1.2 using the machinery of exponential approximations. This technique requires us to couple the original process and the approximations in a way such that the probability of a

non-negligible deviation decays at super-exponential speed. In the present section, we provide details on the coupling construction and show how it can be used to derive Propositions 4.1, 4.2 and 4.3.

5.1. *Total variation bounds for frustrated users.* Since the technique of exponential approximation hinges upon total-variation bounds, it is essential to understand the regularity of frustrated users as a function of the input process. The following result shows that the self-regulating feature of the defining ODE gives rise to excellent continuity properties of the solutions.

We observe that the construction of the process Γ^λ from the Poisson point process Z^λ does not make use of the choice components associated with the Poisson points. Indeed, given the knowledge about the precise locations of the chosen relays, there is no longer any uncertainty on the evolution of frustrated transmitters. Hence, more generally, to any finite counting measure ν on $[0, t_\mathsf{f}] \times W \times Y^\lambda$ we can associate a process of frustrated users $\gamma(\nu)$.

For instance, let $Z^{\lambda,\delta}(\nu_\mathsf{R})$ denote the Poisson point process $[0, t_\mathsf{f}] \times W \times Y^\lambda$ with intensity measure $\lambda\mu^\delta(\nu_\mathsf{R})$ whose density w.r.t. $\mu_\mathsf{T}^\mathsf{t} \otimes \mu_\mathsf{T}^\mathsf{s} \otimes l_\lambda$ is given by

$$\kappa_{\nu_\mathsf{R}}^\delta(y|x) = \sum_{i=1}^k \mathbb{1}\{y \in W_i\} \frac{\kappa_{\nu_\mathsf{R}}(W_i|x)}{l_\lambda(W_i)}.$$

Then, $\gamma(\lambda^{-1}Z^{\lambda,\delta}(\nu_\mathsf{R}))$ coincides in distribution with $\gamma(\mathfrak{L}_\lambda^\delta(\nu_\mathsf{R}), l_\lambda)$. Indeed, having the identity (4.1) at our disposal, we can decompose into the spatial subdomains W_i and then apply Proposition 2.1 in each of these domains separately.

In the following result, we show that γ is 2-Lipschitz on counting measures.

Lemma 5.1. *Let ν, ν' be finite simple counting measures on $[0, t_\mathsf{f}] \times W \times Y^\lambda$. Then,*

$$\|\gamma(\nu) - \gamma(\nu')\| \le 2\|\nu - \nu'\|.$$

Proof: In the proof, we identify ν and ν' with their support and write $\nu^\cup$ and $\nu^\cap$ for their union and intersection, respectively. Then, by monotonicity,

$$\gamma(\nu^\cap) \le \min\{\gamma(\nu), \gamma(\nu')\} \le \max\{\gamma(\nu), \gamma(\nu')\} \le \gamma(\nu^\cup).$$

Hence, it suffices to show $\|\gamma(\nu^\cup) - \gamma(\nu^\cap)\| \le \|\nu^\cup - \nu^\cap\|$. We show this if $\nu^\cup \setminus \nu^\cap$ consists of a singleton $z_0 = \{(t_0, x_0, y_0)\}$. The general statement is obtained via induction. In fact, we can describe precisely how the space-time counting measures $\gamma(\nu^\cup)$ and $\gamma(\nu^\cap)$ differ from each other. If y_0 has already been occupied at time t_0, then $\gamma(\nu^\cup)$ and $\gamma(\nu^\cap)$ agree apart from an additional atom at z_0. On the other hand, if y_0 has not already been occupied at time t_0, then let $z_1 = (t_1, x_1, y_1)$ denote the first particle after time t_0 that points to y_0. If such a particle does not exist, we leave z_1 undefined. Again, $\gamma(\nu^\cup)$ and $\gamma(\nu^\cap)$ agree apart from at most one atom, namely z_1 if defined. $\square$

5.2. *Mixed exponential equivalence.* Next, we consider intermediate approximations which partially replace the limiting relay measure by the empirical measure. For this, we introduce the mixed augmented empirical measures

$$\mathfrak{L}_\lambda^\delta(\mu_\mathsf{R}, l_\lambda) = \sum_{i=1}^k L_\lambda^i(\mu_\mathsf{R}) \otimes 1_{W_i} \frac{l_\lambda}{l_\lambda(W_i)}.$$

Proposition 5.2. *The families of measure-valued processes* $\gamma(\mathfrak{L}_\lambda^\delta(\mu_{\mathsf{R}}, l_\lambda), l_\lambda)$ *and* $\gamma(\mathfrak{L}_\lambda^\delta(\mu_{\mathsf{R}}), \mu_{\mathsf{R}})$ *are* $\|\cdot\|$*-exponentially equivalent.*

Proof: We use the identification (4.1) to decompose $\gamma(\mathfrak{L}_\lambda^\delta(\mu_{\mathsf{R}}, l_\lambda), l_\lambda)$ and $\gamma(\mathfrak{L}_\lambda^\delta(\mu_{\mathsf{R}}), \mu_{\mathsf{R}})$ as

$$\sum_{i=1}^k \gamma\big(\mathfrak{L}_\lambda^\delta(\mu_{\mathsf{R}}, l_\lambda)(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, W_i), l_\lambda(W_i)\big) \text{ and}$$

$$\sum_{i=1}^k \gamma\big(\mathfrak{L}_\lambda^\delta(\mu_{\mathsf{R}})(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, W_i), \mu_{\mathsf{R}}(W_i)\big),$$

respectively. Thus, it suffices to show the exponential equivalence for fixed i. By definition, $\mathfrak{L}_\lambda^\delta(\mu_{\mathsf{R}}, l_\lambda)(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, W_i)$ and $\mathfrak{L}_\lambda^\delta(\mu_{\mathsf{R}})(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, W_i)$ are both empirical measures associated with a Poisson point processes with intensity measure $\lambda \kappa_{\mu_{\mathsf{R}}}(W_i|x)\mu_{\mathsf{T}}(\mathrm{d}s, \mathrm{d}u, \mathrm{d}x)$, so that an application of Lemma 3.8 concludes the proof. $\square$

5.3. *Coupling construction.* The coupling construction announced in the beginning of this section is based on an expansion of the state space $[0, t_{\mathsf{f}}] \times W \times Y^\lambda$ to $V^{*,\lambda} = [0, t_{\mathsf{f}}] \times W \times Y^\lambda \times \mathbb{R}_{\geq 0}$. This allows Poisson point processes of various inhomogeneous intensities to be coupled by considering points whose last coordinate lies below a threshold function. More precisely, let $Z^{\lambda,*}$ denote a Poisson point process on

$$V^{*,\lambda} = [0, t_{\mathsf{f}}] \times W \times Y^\lambda \times \mathbb{R}_{\geq 0}.$$

with intensity measure $\lambda \mu_{\mathsf{T}}^{\mathsf{t}} \otimes \mu_\lambda \otimes |\cdot|$, where $|\cdot|$ is the Lebesgue measure and

$$\mu_\lambda = \mu_{\mathsf{T}}^{\mathsf{s}} \otimes l_\lambda.$$

For any family of measurable functions $f_\lambda : W \times Y^\lambda \to [0, \infty)$ let

$$M(f) = \{(x, y, v) : v \leq f(x, y)\}$$

denote the sub-level set of f. Then, projecting the intersection of $Z^{\lambda,*}$ with $M(f)$ onto $[0, t_{\mathsf{f}}] \times W \times Y^\lambda$ yields a Poisson point process $Z^\lambda(f)$ on $[0, t_{\mathsf{f}}] \times W \times Y^\lambda$ whose intensity measure $\lambda \mu_\lambda^f$ is characterized by

$$\frac{\mathrm{d}\mu_\lambda^f}{\mathrm{d}(\mu_{\mathsf{T}}^{\mathsf{t}} \otimes \mu_\lambda)} = f.$$

For instance, the processes Z^λ and $Z^{\lambda,\delta}(\nu_{\mathsf{R}})$ can be recovered by choosing the threshold function to be $\kappa_{l_\lambda}(y|x)$ and $\kappa_{\nu_{\mathsf{R}}}^\delta(y|x)$, respectively.

For bounded measurable $f, g : W \times Y^\lambda \to [0, \infty)$ we note that the signed counting measure $Z^\lambda(f) - Z^\lambda(g)$ can be decomposed as $Z^{\lambda,+}(f, g) - Z^{\lambda,-}(f, g)$, where

$$Z^{\lambda,+}(f, g) = \{(T_i, X_i, Y_i, V_i) \in Z^{\lambda,*} : g(X_i, Y_i) \leq V_i \leq f(X_i, Y_i)\}$$

and

$$Z^{\lambda,-}(f, g) = \{(T_i, X_i, Y_i, V_i) \in Z^{\lambda,*} : f(X_i, Y_i) \leq V_i \leq g(X_i, Y_i)\}.$$

In particular, the $\|\cdot\|$ distance between $Z^\lambda(f)$ and $Z^\lambda(g)$ can be represented as

$$\|Z^\lambda(f) - Z^\lambda(g)\| = \max\{Z^{\lambda,+}(f, g)(V^{*,\lambda}), Z^{\lambda,-}(f, g)(V^{*,\lambda})\}. \tag{5.1}$$

Thus, for arbitrary $f_\lambda, g_\lambda : W \times Y^\lambda \to [0, \infty)$ the distance $\|Z^\lambda(f) - Z^\lambda(g)\|$ is stochastically bounded by a Poisson random variable with intensity

$$\lambda|f_\lambda - g_\lambda|_{\mu_\lambda} = \lambda \int |f_\lambda(x, y) - g_\lambda(x, y)| \mu_\lambda(\mathrm{d}x, \mathrm{d}y).$$

5.4. *Proof of Propositions 4.1 and 4.2.* We note that Lemma 5.1 and identity (5.1) allow us to reduce the proof of Propositions 4.1 and 4.2 to an intensity bound.

Corollary 5.3. *Let $\{f_\lambda\}_{\lambda>0}$ and $\{f_\lambda^\delta\}_{\delta,\lambda>0}$ denote families of non-negative $L^1(\mu_\lambda)$ functions satisfying*

$$\limsup_{\delta\downarrow 0} \limsup_{\lambda\uparrow\infty} |f_\lambda - f_\lambda^\delta|_{\mu_\lambda} = 0.$$

Then $\gamma(Z^\lambda(f_\lambda^\delta))$ are $\|\cdot\|$-exponentially good approximations of $\gamma(Z^\lambda(f_\lambda))$.

Proof: Let $\varepsilon > 0$ be arbitrary. First, by Lemma 5.1, it suffices to provide suitable bounds on $\mathbb{P}(\|Z^\lambda(f_\lambda) - Z^\lambda(f_\lambda^\delta)\| > \varepsilon)$. Now, the identification (5.1) transforms the distance between $Z^\lambda(f_\lambda)$ and $Z^\lambda(f_\lambda^\delta)$ into the mass of the coupling Poisson process in domains of vanishing μ_λ-measure. Hence, an application of part 3 of Lemma 3.1 concludes the proof. □

Next, we provide an example of an intensity bound that will also be relevant for the identification of the rate function in the following section. For this purpose, we introduce the mixed preference functions

$$\kappa^\delta_{\nu_\mathsf{R}, \nu_\mathsf{R}'}(y|x) = \sum_{i=1}^k \mathbb{1}\{y \in W_i\} \frac{\kappa_{\nu_\mathsf{R}}(W_i|x)}{\nu_\mathsf{R}'(W_i)}$$

and the associated intensity measure $\mu^\delta(\nu_\mathsf{R}, \nu_\mathsf{R}')$ determined by

$$\frac{\mathrm{d}\mu^\delta(\nu_\mathsf{R}, \nu_\mathsf{R}')}{\mathrm{d}(\mu_\mathsf{T}^\mathsf{t} \otimes \mu_\mathsf{T}^\mathsf{s} \otimes \nu_\mathsf{R}')}(t, x, y) = \kappa^\delta_{\nu_\mathsf{R}, \nu_\mathsf{R}'}(y|x).$$

Lemma 5.4. *It holds that $\lim_{\delta\downarrow 0} |\kappa_{\mu_\mathsf{R}} - \kappa^\delta_{\mu_\mathsf{R}, \mu_\mathsf{R}}|_{\mu_\mathsf{T}^\mathsf{s} \otimes \mu_\mathsf{R}} = 0$.*

Proof: Expanding the definitions, we see that the claim is equivalent to proving that

$$\lim_{\delta\downarrow 0} \int_{W^2} \sum_{i=1}^k \mathbb{1}\{y \in W_i\} \Big|\kappa_{\mu_\mathsf{R}}(y|x) - \frac{\int_{W_i} \kappa_{\mu_\mathsf{R}}(y'|x)\mu_\mathsf{R}(\mathrm{d}y')}{\mu_\mathsf{R}(W_i)}\Big| (\mu_\mathsf{T}^\mathsf{s} \otimes \mu_\mathsf{R})(\mathrm{d}x, \mathrm{d}y) = 0.$$

By dominated convergence it suffices to show that the integrand tends to zero for $\mu_\mathsf{T}^\mathsf{s} \otimes \mu_\mathsf{R}$-almost every (x, y). But this is a consequence of the Lebesgue Density Theorem (Rudin, 1987). □

Now, we can prove Propositions 4.1 and 4.2.

Proof of Propositions 4.1 and 4.2: By Proposition 5.2, Corollary 5.3 and Lemma 5.4 it suffices to show that

$$\lim_{\lambda\uparrow\infty} |\kappa^\delta_{\mu_\mathsf{R}, l_\lambda} - \kappa^\delta_{l_\lambda, l_\lambda}|_{\mu_\lambda} = 0 \tag{5.2}$$

for every $\delta > 0$, and

$$\lim_{\lambda\uparrow\infty} |\kappa_{l_\lambda} - \kappa^\delta_{l_\lambda, l_\lambda}|_{\mu_\lambda} = |\kappa_{\mu_\mathsf{R}} - \kappa^\delta_{\mu_\mathsf{R}, \mu_\mathsf{R}}|_{\mu_\mathsf{T}^\mathsf{s} \otimes \mu_\mathsf{R}}. \tag{5.3}$$

We begin by considering (5.2). First, $|\kappa^{\delta}_{\mu_{\mathsf{R}}, l_{\lambda}} - \kappa^{\delta}_{l_{\lambda}, l_{\lambda}}|_{\mu_{\lambda}}$ is given by

$$\sum_{i=1}^{k} \int_{W} |\kappa_{\mu_{\mathsf{R}}}(W_i|x) - \kappa_{l_{\lambda}}(W_i|x)| \mu^{\mathsf{s}}_{\mathsf{T}}(\mathrm{d}x).$$

In particular, by dominated convergence, it suffices to show that the integrand converges to zero for $\mu^{\mathsf{t}}_{\mathsf{T}}$-almost every $x \in W$. Hence, let $1 \le i \le k$ and $x \in W$ be arbitrary. Then, $\kappa_{\mu_{\mathsf{R}}}(W_i|x) - \kappa_{l_{\lambda}}(W_i|x)$ is given by

$$\frac{\int_{W_i} \kappa(x, y) \mu_{\mathsf{R}}(\mathrm{d}y)}{\int_{W} \kappa(x, y) \mu_{\mathsf{R}}(\mathrm{d}y)} - \frac{\int_{W_i} \kappa(x, y) l_{\lambda}(\mathrm{d}y)}{\int_{W} \kappa(x, y) l_{\lambda}(\mathrm{d}y)}.$$

Disregarding a $\mu^{\mathsf{t}}_{\mathsf{T}}$-nullset, we may assume that $\kappa(x, \cdot)$ is μ_{R}-almost everywhere continuous, so that the weak convergence $l_{\lambda} \to \mu_{\mathsf{R}}$ implies (5.2).

For the proof of (5.3) note that, by dominated convergence, it suffices to show that for $\mu^{\mathsf{s}}_{\mathsf{T}}$-almost every x,

$$\lim_{\lambda \uparrow \infty} |\kappa_{l_{\lambda}}(\cdot|x) - \kappa^{\delta}_{l_{\lambda}, l_{\lambda}}(\cdot|x)|_{l_{\lambda}} = |\kappa_{\mu_{\mathsf{R}}}(\cdot|x) - \kappa^{\delta}_{\mu_{\mathsf{R}}, \mu_{\mathsf{R}}}(\cdot|x)|_{\mu_{\mathsf{R}}}.$$

First, as μ_{R} is the weak limit of l_{λ}, both $|\kappa_{l_{\lambda}}(\cdot|x) - \kappa_{\mu_{\mathsf{R}}}(\cdot|x)|_{l_{\lambda}}$ and $|\kappa^{\delta}_{l_{\lambda}, l_{\lambda}}(\cdot|x) - \kappa^{\delta}_{\mu_{\mathsf{R}}, \mu_{\mathsf{R}}}(\cdot|x)|_{l_{\lambda}}$ tend to zero as λ tends to infinity. Therefore, it remains to show that

$$\lim_{\lambda \uparrow \infty} |\kappa_{\mu_{\mathsf{R}}}(\cdot|x) - \kappa^{\delta}_{\mu_{\mathsf{R}}, \mu_{\mathsf{R}}}(\cdot|x)|_{l_{\lambda}} = |\kappa_{\mu_{\mathsf{R}}}(\cdot|x) - \kappa^{\delta}_{\mu_{\mathsf{R}}}(\cdot|x)|_{\mu_{\mathsf{R}}},$$

which again is a consequence of the weak convergence of the relay measure. $\qquad\square$

5.5. *Coupling construction for absolutely continuous measures.* It should not come as a surprise that similar to what we have seen in the empirical setting in Section 5.3, couplings play a vital rôle in the identification of the rate function. The procedure in the absolutely continuous setting is very similar, but some care has to be taken since the empirical measures l_{λ} need to be replaced by the limiting measure μ_{R} and the unit interval is added to the state space. More precisely, we consider measures on the space

$$V^* = [0, t_{\mathsf{f}}] \times [0, 1] \times W^2 \times [0, \kappa_{\infty}],$$

where $\kappa_{\infty} = \sup_{x, y \in W} \kappa_{\mu_{\mathsf{R}}}(y|x)$. To simplify notation, we write κ and $\gamma(\cdot)$ instead of the more verbose $\kappa_{\mu_{\mathsf{R}}}$ and $\gamma(\cdot, \mu_{\mathsf{R}})$. If $f : W^2 \to [0, \kappa_{\infty}]$ is a measurable function and $\mathfrak{n}^* \in \mathcal{M}^* = \mathcal{M}(V^*)$, then we let $\mathfrak{n}^*(f)$ denote the measure on V^* that is defined by restriction to the sublevel set $M(f)$ and forgetting the last coordinate. For instance, we can recover previously introduced intensity measures as

$$\mu(\mu_{\mathsf{R}}) = \mu^*(\kappa) \qquad \text{and} \qquad \mu^{\delta}(\mu_{\mathsf{R}}, \mu_{\mathsf{R}}) = \mu^*(\kappa^{\delta}),$$

where

$$\mu^* = \mu^{\mathsf{t}}_{\mathsf{T}} \otimes U \otimes \mu^{\mathsf{s}}_{\mathsf{T}} \otimes \mu_{\mathsf{R}} \otimes |\cdot|.$$

The decisive advantage offered by the couplings is that they allow for an elegant way of expressing total-variation distances. More precisely, for bounded measurable $f, g : W^2 \to [0, \kappa_{\infty}]$ we note that the signed measure $\mathfrak{n}^*(f) - \mathfrak{n}^*(g)$ can be decomposed as $\mathfrak{n}^{*,+}(f, g) - \mathfrak{n}^{*,-}(f, g)$, where

$$\frac{\mathrm{d}\mathfrak{n}^{*,+}(f, g)}{\mathrm{d}\mathfrak{n}^*}(t, u, x, y, v) = \mathbb{1}\{g(x, y) \le v \le f(x, y)\}$$

and
$$\frac{\mathrm{d}\mathfrak{n}^{*,-}(f,g)}{\mathrm{d}\mathfrak{n}^*}(t,u,x,y,v) = \mathbb{1}\{f(x,y) \le v \le g(x,y)\}.$$
In particular,
$$\|\mathfrak{n}^*(f) - \mathfrak{n}^*(g)\| = \max\{\mathfrak{n}^{*,+}(f,g)(V^*), \mathfrak{n}^{*,-}(f,g)(V^*)\}. \tag{5.4}$$

5.6. *Uniform spatial approximation property.* The coupling introduced in the previous section brings us into the setting of Dembo and Zeitouni (1998, Theorem 4.2.23) where both the desired rate functions and the rate functions of the approximations are of contraction type. For the rate-function approximation to be useful in the exponential-approximation argument, we need to verify that the approximations are uniform on measures with bounded entropy, i.e., on
$$I_\alpha = \{\mathfrak{n}^* \in \mathcal{M}^* : h(\mathfrak{n}^*|\mu^*) \le \alpha\}.$$

Lemma 5.5. *Let $\alpha > 0$ be arbitrary. Then,* $\lim_{\delta \downarrow 0} \sup_{\mathfrak{n}^* \in I_\alpha} \|\gamma(\mathfrak{n}^*(\kappa)) - \gamma(\mathfrak{n}^*(\kappa^\delta))\|$ *$= 0$.*

Before we prove Lemma 5.5, we explain how it can be used to derive Proposition 4.3.

Proof of Proposition 4.3: Although Lemma 5.5 is the main ingredient for the exponential approximation (Dembo and Zeitouni, 1998, Theorem 4.2.23), there are still two further steps that remain to be verified. First, we need to check that the contraction-type rate functions
$$\inf_{\mathfrak{n}^* \in \mathcal{M}^*:\, (\gamma_{t_i}(\mathfrak{n}^*(\kappa)))_i = (\gamma_{t_i})_i} h(\mathfrak{n}^*|\mu^*) \qquad \text{and} \qquad \inf_{\mathfrak{n} \in \mathcal{M}':\, (\gamma_{t_i}(\mathfrak{n}))_i = (\gamma_{t_i})_i} h(\mathfrak{n}|\mu(\mu_\mathsf{R}))$$
are identical. Second, the continuity of the map $I_\alpha \to \mathcal{M}(W)^{[0,t_\mathsf{f}]}$, $\mathfrak{n}^* \mapsto \gamma(\mathfrak{n}^*(\kappa^\delta))$ needs to be justified. In order to verify the identity of the rate functions, we prove that
$$\inf_{\mathfrak{n}^* \in \mathcal{M}^*:\, \mathfrak{n}^*(\kappa) = \mathfrak{n}} h(\mathfrak{n}^*|\mu^*) = h(\mathfrak{n}|\mu(\mu_\mathsf{R})). \tag{5.5}$$
Showing that the l.h.s. is at most as large as the r.h.s. is achieved by setting
$$\mathfrak{n}_0^*(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, \mathrm{d}y, \mathrm{d}v) = \kappa(y|x)^{-1}\mathbb{1}\{M(\kappa)\}\mathfrak{n}(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, \mathrm{d}y)\mathrm{d}v$$
$$+ \mathbb{1}\{M(\kappa)^c\}\mu^*(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, \mathrm{d}y, \mathrm{d}v).$$
For the reverse inequality it can be checked by direct computation that
$$h(\mathfrak{n}^*|\mu^*) = h(\mathfrak{n}^*|\mathfrak{n}_0^*) + h(\mathfrak{n}|\mu(\mu_\mathsf{R})),$$
so that the non-negativity gives (5.5).

Second, we show that for fixed $\delta > 0$ the map $\mathfrak{n}^* \mapsto \gamma(\mathfrak{n}^*(\kappa^\delta))$ is continuous. For this note that $\gamma(\mathfrak{n}^*(\kappa^\delta))$ decomposes as
$$\gamma(\mathfrak{n}^*(\kappa^\delta)) = \sum_{i=1}^{k} \gamma(\mathfrak{n}_i^*(\kappa^\delta)),$$
where
$$\mathfrak{n}_i^*(\kappa^\delta)(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x) = \mathfrak{n}^*(\kappa^\delta)(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, W_i)$$
is the restriction of $\mathfrak{n}^*(\kappa^\delta)$ to the points whose y-coordinate is in W_i. Similarly, put
$$\mu_i(\mu_\mathsf{R})(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x) = \mu^\delta(\mu_\mathsf{R}, \mu_\mathsf{R})(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, W_i).$$

Since $h(\mathfrak{n}_i^*(\kappa^\delta)|\mu_i(\mu_\mathsf{R})) \leq \alpha$, we deduce from Proposition 2.4 that γ is continuous at $\mathfrak{n}_i^*(\kappa^\delta)$. Combining this observation with the continuity of the partial evaluation maps $\mathfrak{n}^* \mapsto \mathfrak{n}_i^*(\kappa^\delta)$ concludes the proof. $\qquad\square$

Hence, it remains to prove Lemma 5.5. We recall from equation (1.9) that the process $\beta(\mathfrak{n})$ is obtained as a spatial mixture of the corresponding localized processes $\beta(\mathfrak{n}_y)$ at receiver locations $y \in W$. Therefore, understanding how sensitive the localized processes are w.r.t. their input measures lays the groundwork for the global setting.

Lemma 5.6. *Let $\nu, \nu' \in \mathcal{M}_\mathsf{ac}$, then $\|\beta(\nu)(W) - \beta(\nu')(W)\| \leq \|\nu - \nu'\|$.*

Proof: Let $t \in [0, t_\mathsf{f}]$ be arbitrary. By symmetry, it suffices to derive an upper bound for $\beta_t(\nu) - \beta_t(\nu')$, where for ease of notation we omit the evaluation on W. Now, let $t_0 \in [0, t_\mathsf{f}]$ be the last point before t such that

$$\beta_{t_0}(\nu) \leq \beta_{t_0}(\nu'),$$

then,

$$\beta_t(\nu) - \beta_t(\nu') = \int_{t_0}^t \nu(\mathrm{d}s, [\beta_s(\nu), 1], W) - \int_{t_0}^t \nu'(\mathrm{d}s, [\beta_s(\nu'), 1], W).$$

This difference can be split up into

$$\int_{t_0}^t \nu(\mathrm{d}s, [\beta_s(\nu), 1], W) - \int_{t_0}^t \nu(\mathrm{d}s, [\beta_s(\nu'), 1], W)$$

and

$$\int_{t_0}^t \nu(\mathrm{d}s, [\beta_s(\nu'), 1], W) - \int_{t_0}^t \nu'(\mathrm{d}s, [\beta_s(\nu'), 1], W),$$

where we know that the first expression is negative and therefore can be omitted. Observing that the second expression is at most $\|\nu - \nu'\|$ concludes the proof. $\quad\square$

Now we use Lemma 5.6 to complete the derivation of Lemma 5.5.

Proof of Lemma 5.5: First,

$$\|\mathfrak{n}^*(\kappa)([0, t], [0, 1], \mathrm{d}x) - \mathfrak{n}^*(\kappa^\delta)([0, t], [0, 1], \mathrm{d}x)\| \leq \|\mathfrak{n}^*(\kappa) - \mathfrak{n}^*(\kappa^\delta)\|$$

so that by Lemma 3.1 part 2, Lemma 5.4 and identity (5.4) it remains to prove the statement with γ replaced by β. By absolute continuity, we can perform disintegration of the measures $\mathfrak{n}^*(\kappa)$ and $\mathfrak{n}^*(\kappa^\delta)$ with respect to the relay coordinate. That is,

$$\mathfrak{n}^*(\kappa)(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, \mathrm{d}y) = \mathfrak{n}_y^*(\kappa)(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x)\mu_\mathsf{R}(\mathrm{d}y)$$

and

$$\mathfrak{n}^*(\kappa^\delta)(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x, \mathrm{d}y) = \mathfrak{n}_y^*(\kappa^\delta)(\mathrm{d}t, \mathrm{d}u, \mathrm{d}x)\mu_\mathsf{R}(\mathrm{d}y).$$

Let $t \in [0, t_\mathsf{f}]$ and $A \subset W$ measurable. Then, inserting the definition of β we see that we need to compare

$$\int_{[0,t] \times W} \mathfrak{n}_y^*(\kappa)(\mathrm{d}s, [\beta_s(\mathfrak{n}_y^*(\kappa))(W), 1] \times A)\mu_\mathsf{R}(\mathrm{d}y)$$

with

$$\int_{[0,t] \times W} \mathfrak{n}_y^*(\kappa^\delta)(\mathrm{d}s, [\beta_s(\mathfrak{n}_y^*(\kappa^\delta))(W), 1] \times A)\mu_\mathsf{R}(\mathrm{d}y).$$

We decompose this task into providing bounds separately for

$$\left| \int_{[0,t]\times W} (\mathfrak{n}_y^*(\kappa) - \mathfrak{n}_y^*(\kappa^\delta))(\mathrm{d}s, [\beta_s(\mathfrak{n}_y^*(\kappa^\delta))(W), 1] \times A)\mu_\mathsf{R}(\mathrm{d}y) \right|$$

and

$$\int_{[0,t_\mathsf{f}]\times W} \mathfrak{n}_y^*(\kappa)(\mathrm{d}s, \mathcal{I}(\beta_s(\mathfrak{n}_y^*(\kappa))(W), \beta_s(\mathfrak{n}_y^*(\kappa^\delta))(W)) \times W)\mu_\mathsf{R}(\mathrm{d}y),$$

where $\mathcal{I}(a,b) = [a \wedge b, a \vee b]$. The first expression is at most $\|\mathfrak{n}^*(\kappa) - \mathfrak{n}^*(\kappa^\delta)\|$, so that again Lemma 3.1 part 2, Lemma 5.4 and identity (5.4) yield that

$$\lim_{\delta\downarrow 0} \sup_{\mathfrak{n}^*\in I_\alpha} \|\mathfrak{n}^*(\kappa) - \mathfrak{n}^*(\kappa^\delta)\| = 0.$$

By Lemma 5.6, the second expression is bounded above by $\mathfrak{n}^*(\kappa)(C_{\mathfrak{n}^*,\delta})$ where

$$C_{\mathfrak{n}^*,\delta} = \{(t,u,x,y): |u - \beta_t(\mathfrak{n}_y^*(\kappa))(W)| \leq \|\mathfrak{n}_y^*(\kappa) - \mathfrak{n}_y^*(\kappa^\delta)\|\}.$$

In particular, by Lemma 3.1 part 2 it remains to show that $\lim_{\delta\downarrow 0}\sup_{\mathfrak{n}^*\in I_\alpha} \mu(\mu_\mathsf{R})(C_{\mathfrak{n}^*,\delta}) = 0$. For this, we note that reversing the disintegration of the relay measure gives that

$$\mu(\mu_\mathsf{R})(C_{\mathfrak{n}^*,\delta}) \leq 2\int_{W^2} \|\mathfrak{n}_y^*(\kappa) - \mathfrak{n}_y^*(\kappa^\delta)\|\kappa(y|x)(\mu_\mathsf{T}^\mathsf{s} \otimes \mu_\mathsf{R})(\mathrm{d}x, \mathrm{d}y)$$

$$\leq 2\mu_\mathsf{T}^\mathsf{s}(W)\kappa_\infty \int_W \max\{\mathfrak{n}_y^{*,+}(\kappa, \kappa^\delta), \mathfrak{n}_y^{*,-}(\kappa, \kappa^\delta)\}\mu_\mathsf{R}(\mathrm{d}y)$$

$$\leq 2\mu_\mathsf{T}^\mathsf{s}(W)\kappa_\infty(\mathfrak{n}^{*,+}(\kappa, \kappa^\delta) + \mathfrak{n}^{*,-}(\kappa, \kappa^\delta)),$$

so that another invocation of Lemma 3.1 part 2 and Lemma 5.4 concludes the proof. □

Acknowledgments

References

D. Bethanabhotla, O. Y. Bursalioglu, H. C. Papadopoulos and G. Caire. Optimal user-cell association for massive MIMO wireless networks. *IEEE Transactions on Wireless Communications* **15** (3), 1835–1850 (2016). DOI: 10.1109/TWC.2015.2496942.

S. Boucheron, F. Gamboa and C. Léonard. Bins and balls: large deviations of the empirical occupancy process. *Ann. Appl. Probab.* **12** (2), 607–636 (2002). MR1910642.

S. Boucheron, G. Lugosi and P. Massart. *Concentration inequalities.* Oxford University Press, Oxford (2013). ISBN 978-0-19-953525-5. MR3185193.

D. J. Daley and D. Vere-Jones. *An introduction to the theory of point processes.* Springer Series in Statistics. Springer-Verlag, New York (1988). ISBN 0-387-96666-8. MR950166.

A. Dembo and O. Zeitouni. *Large deviations techniques and applications*, volume 38 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, second edition (1998). ISBN 0-387-98406-2. MR1619036.

P. Eichelsbacher and U. Schmock. Exponential approximations in completely regular topological spaces and extensions of Sanov's theorem. *Stochastic Process. Appl.* **77** (2), 233–251 (1998). MR1649006.

J. Feng and T. G. Kurtz. *Large deviations for stochastic processes*, volume 131 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI (2006). ISBN 978-0-8218-4145-7; 0-8218-4145-9. MR2260560.

C. Graham and S. Méléard. A large deviation principle for a large star-shaped loss network with links of capacity one. *Markov Process. Related Fields* **3** (4), 475–492 (1997a). MR1607107.

C. Graham and S. Méléard. An upper bound of large deviations for a generalized star-shaped loss network. *Markov Process. Related Fields* **3** (2), 199–223 (1997b). MR1468174.

A. Granas and J. Dugundji. *Fixed point theory*. Springer Monographs in Mathematics. Springer-Verlag, New York (2003). ISBN 0-387-00173-5. MR1987179.

C. Léonard. Large deviations for long range interacting particle systems with jumps. *Ann. Inst. H. Poincaré Probab. Statist.* **31** (2), 289–323 (1995). MR1324810.

W. Rudin. *Real and complex analysis*. McGraw-Hill Book Co., New York, third edition (1987). ISBN 0-07-054234-1. MR924157.

A. Sankararaman and F. Baccelli. Spatial birth-death wireless networks. *IEEE Trans. Inform. Theory* **63** (6), 3964–3982 (2017). MR3677758.

A. Shwartz and A. Weiss. *Large deviations for performance analysis*. Stochastic Modeling Series. Chapman & Hall, London (1995). ISBN 0-412-06311-5. MR1335456.

A. Shwartz and A. Weiss. Large deviations with diminishing rates. *Math. Oper. Res.* **30** (2), 281–310 (2005). MR2142034.