A Performance Analysis Study of Multipath Routing in a Hybrid Network with Mobile Users

Andrei Iu. Bejan*, Richard J. Gibbens*, Yeon-sup Lim†, Don Towsley†
*Computer Laboratory, University of Cambridge, 15 JJ Thomson Avenue, Cambridge CB23 8SP, UK.
†Department of Computer Science, University of Massachusetts, Amherst MA 01003 USA.
Emails: andrei.bejan@cl.cam.ac.uk, richard.gibbens@cl.cam.ac.uk, ylim@cs.umass.edu, towsley@cs.umass.edu

***Abstract*—Mobile communication platforms of individual agents and ground vehicles, ships and aircrafts of both civil and military services often operate in highly dynamic conditions with constantly changing infrastructure and access to communication resources. Efficient techniques for rapid and yet stable communication of such fleets with their control centres and between cooperating vehicles within the fleet is a challenging but important area of study with the potential to facilitate the analysis and design of efficient and robust communication systems. Multipath extensions of data transmission protocols aim to take advantage of path diversity to achieve efficient and robust bandwidth allocation while maintaining stability. Such multipath resource pooling extensions of routing and congestion control intrinsically implement decentralisation with implicit resource sharing. In this paper, we build on the recent theoretical work on fluid model approximations of multipath TCP and study their application to the scenarios in which a convoy with two communication nodes (representing the convoy's head and tail) establishes channels with a set of radio/WiFi towers and a satellite relaying information to a remote destination; these channels have time-varying capacities which depend on the position and dynamics of the convoy. The paper studies the performance of a multipath TCP controller and demonstrates how path diversity can be implicitly utilised to spread flows across available paths. Furthermore, we study the patterns of sub-flows governed by dynamic control according to the motion of the convoy and investigate the trade-offs between resource utilisation and the speed of response by the sub-flows.**

I. Introduction

Public and military services often operate in dynamic environments which offer constantly changing infrastructure and access to communication resources. In particular, mobile communication platforms of ground vehicles (including taxi, emergency and public transport fleets), ships and aircrafts, experience fluctuating or disrupting radio links, bandwidths and connectivity. Efficient techniques for rapid and yet stable communication of such fleets with their control centres and between cooperating vehicles within the fleet is a challenging but important area of study.

Communication based on single end-to-end paths could easily expose data transmission, and, in particular continuous data streaming to the risk of instability and disruption. Multipath extensions of data transmission protocols aim to take advantage of path diversity to achieve efficient bandwidth allocation while maintaining stability and connectivity. Such multipath resource pooling extensions of data routing and congestion control implement decentralisation with implicit resource sharing and, at the same time, may advantageously implement coordinated control where the rates over available paths are determined as a function of all or some of the available paths [7]. Mobile units may use hybrid networks with potential access to a rich set of technologies and paths (e.g. 3G, 4G, 802.11, satellite). Hence, there is both a necessity and opportunity for robust multipath data transport in such mobile networks.

Practical implementation aspects of congestion control in multipath transport protocols were considered in [14], [16], [21]. Opportunistic mobility with Multipath TCP (MPTCP) for usage of small devices in indoor environments was studied via simulation and indoor WiFi and 3G experiments in [15], where it has been shown that MPTCP gives better throughput, achieves smoother handoffs, and can also be used to lower energy consumption. Also, the feasibility of using MPTCP for mobile/WiFi handover in the current Internet was explored experimentally over real WiFi/3G networks in [11].

Design and stability of multipath end-to-end congestion and rate/routing control algorithms can be studied and conveniently formulated at the scale of flow rates. In recent years researchers have developed a framework that allows congestion control algorithms to be interpreted as distributed mechanisms solving a global optimisation problem [5], [12], [13], [17], [18]. The framework is based on fluid-flow models, and the form of the optimisation problem makes explicit the equilibrium resource allocation policy which can often be restated in terms of a fairness criterion.

In this paper, we build on the recent theoretical work on fluid model approximations of multipath TCP [2], [4], [6], [20] as well as the simulation experiments presented in [9] and further numerical investigations [1]. We study the performance of a scalable multipath TCP controller used in scenarios with the mobility of users. Specifically, we consider the application of fluid models to a setting in which a convoy with two communication nodes (representing a convoy's head and tail) establishes channels with a set of radio/WiFi towers together with a satellite relaying information to a remote destination—these channels have capacities which vary according to the time-dependent position and dynamics of the convoy.

This paper is organised as follows. Section II contains a brief formal analytical description of a fluid model approximation of multipath TCP. This formulation includes a description of the network resources, sources of flow, the available routes

and an algorithm for flow rates control. We then proceed with the description of a convoy scenario with time-varying resource capacities and whose dynamics is described in terms of the aforementioned fluid model controller. Section III presents numerical simulation experiments in which the convoy follows a trajectory with non-constant speed and we explore the behaviour of the resulting dynamic network with time-varying channel bandwidths. Section IV contains further discussion and conclusions.

II. Model Description

A. Fluid model approximation of multipath TCP

We describe a communication network and a routing extension to a TCP algorithm which optimally splits the flows between each source-destination pair and can be implemented at the communication sources. This algorithm is fully decentralised in the sense that there exists a sufficient condition for the local stability of each of the routes which depends on the round-trip times of that route, but not on the round-trip times of other routes. Another feature of this scheme is that the control which a source has over its available route flow rates is treated similarly to link congestion feedback.

1) Network description: We assume that the network consists of a collection of *sources* S and collection of *resources* J. Each source $s \in S$ identifies a unique source-destination pair which is also associated with a set of *routes* R. Each such route $r \in R$ is a set of resources, that is $r \subseteq J$. If a source s transmits along a route r, then we write $r \in s$. Similarly, if a route r uses a resource $j \in J$, we write $j \in r$. Furthermore, the notation $s(r)$ means the unique source identified with the route r.

For each route r there is an associated flow rate $x_r(t) \geq 0$ which represents a dynamic fluid approximation to the rate at which the source $s(r)$ is sending packets along the route r at time t.

Let T_{rj} be the propagation delay from source $s(r)$ to resource j, that is the length of time it takes for a packet to travel from $s(r)$ to j along route r, and let T_{jr} be the propagation delay from resource j to source $s(r)$, that is the length of time it takes for congestion control feedback to reach $s(r)$ from j along r. (It is assumed that a packet must reach its destination before an acknowledgment containing congestion feedback is returned to its source.) The *round trip time* for route r is then given by $T_r = T_{rj} + T_{jr}$ for all $j \in r$.

Finally, we use the notation $a = [\, b \,]_c^+$ defined for $c \geq 0$ to mean $a = b$ if $c > 0$ and $a = \max(0, b)$ if $c = 0$.

2) Rate control algorithm: We use the above description and consider the routing extension to TCP given in [4] briefly summarising its primal version here. (See [6] for further discussion of scalable TCP and [2] for further background on primal multipath algorithms.) In this fluid-flow model of joint routing and rate control the flow rates, $x_r(t)$, vary according to the following differential equations:

$$\dot{x}_r(t) = \frac{x_r(t - T_r)}{T_r} \left[\, \overline{a}(1 - \lambda_r(t)) - b_r y_{s(r)}(t) \lambda_r(t) \,\right]^+_{x_r(t)}, \tag{1}$$

where

$$\lambda_r(t) = 1 - \prod_{j \in r} (1 - \mu_j(t - T_{jr})), \tag{2}$$

$$y_s(t) = \sum_{\tilde{r} \in s} x_{\tilde{r}}(t - T_{\tilde{r}}), \tag{3}$$

and

$$\mu_j(t) = p_j \left(\sum_{\tilde{r}: j \in \tilde{r}} x_{\tilde{r}}(t - T_{\tilde{r}j}) \right). \tag{4}$$

In the above formulation $\overline{a}$, b_r and functions $\lambda_r(\cdot)$ ($r \in R$), $\mu_j(\cdot)$, and $p_j(\cdot)$ ($j \in J$) have the following interpretations. Each resource j has a capacity $c_j > 0$ and exhibits congestion by dropping or marking packets via a penalty function given by

$$p_j(z_j) = \left(\frac{z_j}{c_j} \right)^{\beta_j} \tag{5}$$

for some constant $\beta_j > 0$ (also known as the *resource responsiveness*). Thus, $\mu_j(t)$ is the resource j's dropping/marking rate at time t and $\lambda_r(t)$ is the route r's proportion of the acknowledgements that indicate congestion. The quantity $\overline{a}$ is a proportionality factor of the amount by which the sending rate is increased (on a receipt of a positive acknowledgement), and b_r is a route-specific proportionality factor of the amount by which the sending rate is decreased (on a receipt of a negative acknowledgement through a timeout and so indicating congestion).

Note that equations (1)-(4) with r spanning the whole set of routes R define a set of coupled first-order non-linear differential equations with discrete time delays. Hence, in order to obtain a particular flow dynamics $x_r(t)$, $r \in R$, in a network topology with at least one shared resource one needs to solve the entire system.

3) Local stability condition: A sufficient condition for local stability of the above (primal) algorithm with the penalty function as in (5) is that [4]

$$\overline{a}(1 + \max_{j \in J} \beta_j) < \frac{\pi}{2}. \tag{6}$$

A general form of the sufficient stability condition for arbitrary cost functions p_j is also derived in [4] and is stated in terms of the utility function used for formulating an associated optimisation problem. The choice of the utility function may correspond to α-fairness [5], [10], [18] (e.g. maximum throughput, proportionally fair, or max-min fair allocation scheme). The equilibrium points of the controller described above in equation (6) solve a corresponding optimisation problem to achieve proportionally fair allocation where $\alpha = 1$.

B. Communication network under convoy scenario

1) General description and network topology: Consider a group of vehicles moving in a convoy along a certain path at a possibly variable speed. We shall assume for the purpose of illustration that this fleet moves in an environment in which the only communication channels that can be established to the outside of the fleet are those with a satellite and three

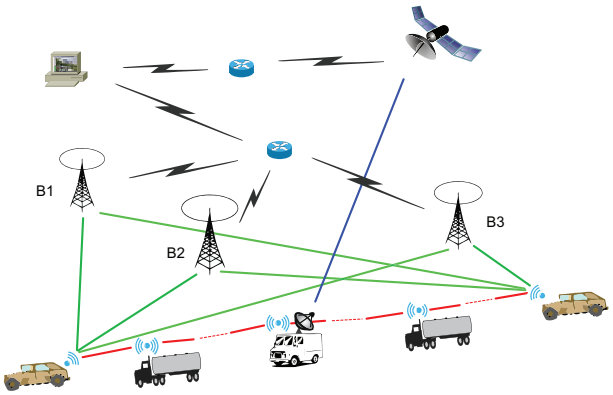

Fig. 1. This figure depicts a moving convoy scenario where both the head and tail of the convoy have access to three WiFi base stations, and the vehicle in the middle of the convoy establishes a connection to a satellite. Information can be transmitted instantaneously within the convoy, thus making it a single communication source transmitting to a destination (i.e. control centre) via multiple paths (see also Figure 3).

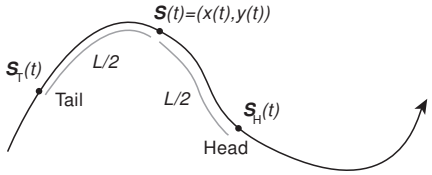

Fig. 2. This figure shows how the convoy's movement is modelled. The centre of the convoy, its head and tail follow a smooth path $\mathbf{S}(t)$, where t is a time parameter, so that the distance between the middle of the convoy to its both head and tail on $\mathbf{S}(t)$ is half the convoy's length.

radio base stations (see Figure 1). Henceforth, we shall use the term 'WiFi' to refer to these base stations and communication channels they provide in an abstract sense of radio communication rather than any specific standard they might use (e.g. IEEE 802.11), as this is not central to our paper. Furthermore, we shall assume that the satellite channel can only be used by a single vehicle situated in the middle or any other part of the fleet (this channel's parameters will remain constant throughout each of our experiments), whereas WiFi channels are only accessible by the fleet's head and tail. We assume that individual vehicles in this group can instantaneously communicate with each other. The whole fleet is considered as a single communication source streaming data to its destination (e.g. a control centre) via available WiFi and satellite channels.

2) Formulation in terms of the fluid model description: As the convoy fleet moves, its head and tail change their positions and this results in time-varying WiFi communication channel capacities. Next, we provide a formal description of this scenario in terms of the model presented in Section II-A.

Let the position of the convoy's centre at time t be described (in the plane) by $\mathbf{S}(t) = (X(t), Y(t))$, where $X(t)$ and $Y(t)$ are some smooth functions used here to describe the convoy's trajectory. We model the movement of the convoy's head and tail as follows: at time t the convoy's head is located at $\mathbf{S}_{\mathrm{H}}(t) = \mathbf{S}(t + \tau_{\mathrm{H}}(t))$ and its tail is located at

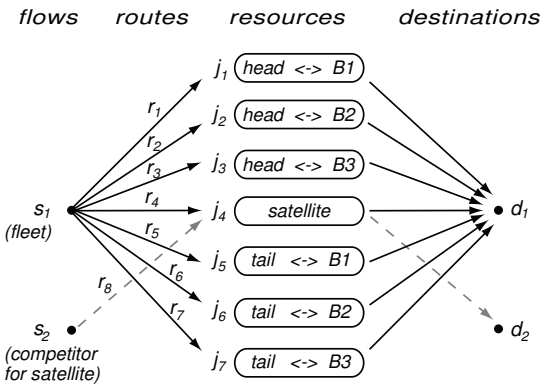

Fig. 3. **Solid lines**: The main example network of one source (convoy fleet) and one destination (i.e. control centre), seven resources (six WiFi channels and one satellite channel) and a total of seven routes. **Solid and dashed lines**: The extended model with a competitor for satellite discussed in Section III-F.

$\mathbf{S}_{\mathrm{T}}(t) = \mathbf{S}(t + \tau_{\mathrm{T}}(t))$, where the functions $\tau_{\mathrm{H}}(t)$ and $\tau_{\mathrm{T}}(t)$ are such that

$$L/2 = \int_{t}^{t+\tau_{\mathrm{H}}(t)} \sqrt{[\dot{X}(\tau)]^2 + [\dot{Y}(\tau)]^2}\, \mathrm{d}\tau, \tag{7}$$

$$L/2 = \int_{t-\tau_{\mathrm{T}}(t)}^{t} \sqrt{[\dot{X}(\tau)]^2 + [\dot{Y}(\tau)]^2}\, \mathrm{d}\tau, \tag{8}$$

and where L is the convoy's length (see Figure 2). For example, if the parametrisation $\mathbf{S}(t)$ is such that $|\dot{\mathbf{S}}(t)|$ does not depend on t, that is the convoy is moving with a constant speed $v = |\dot{\mathbf{S}}(t)|$, then, trivially, $\tau_{\mathrm{H}} = \tau_{\mathrm{T}} = L/2v$.

In most of our experiments the communication network (S, R, J) consists of a single source s_1 representing the convoy, seven resources (six WiFi channels and a single satellite available), and seven routes in total each using a separate resource, as shown in Figure 3 (solid graphics). In a further experiment this model is extended by adding a competitor s_2 who shares the satellite channel, as shown in the same figure with dashed lines (this model is studied in Section III-F).

We assume that the WiFi towers (base stations) are positioned at locations $\mathrm{B}_1 = (-200, -400)$, $\mathrm{B}_2 = (-300, 400)$, $\mathrm{B}_3 = (-100, -200)$ in the plane. As the convoy fleet moves, the actual capacities of the WiFi channels r_1, r_2, r_3, r_5, r_6, r_7 change as follows:

$$c_{j_k}(t) = \rho(d(\mathbf{S}_{\mathrm{H}}(t), \mathrm{B}_k))C_{j_k},\ k = 1, 2, 3, \tag{9}$$

and

$$c_{j_k}(t) = \rho(d(\mathbf{S}_{\mathrm{T}}(t), \mathrm{B}_k))C_{j_k},\ k = 5, 6, 7, \tag{10}$$

where $\rho(d)$ is a decay function measuring the actual channel capacity as a function of distance between a WiFi station and a communicating node, $d(\cdot, \cdot)$ stands for the Euclidean distance between two points, and C_{j_k} is the nominal capacity of the channel/resource j_k. The decay function $\rho(d)$, which we refer to as a *capacity factor*, is assumed to be a non-negative non-increasing function, such that $\rho(0) = 1$. In practical system

deployments channel capacities may vary with and depend on many factors including the terrain and obstacles present; however, in this paper we do not aim to model accurately power attenuation and channel communication schemes (see a discussion on this in e.g. [19])—instead, we make an idealised and yet simple and flexible choice which allows us to take into account dependence of the radio channel's maximal throughput on distance between the tower and mobile agent. We use the straightforward logistic function to represent the capacity factor in our multipath communication model with mobile users as follows:

$$\rho(d) = \frac{[1-(1+e^{-\alpha(d-d_0)})^{-1}]}{[1-(1+e^{\alpha d_0})^{-1}]},\ \alpha > 0, d_0 > 0. \qquad (11)$$

Figure 4 shows two examples of how this decay function can be used to model the dependency of the capacity factor on distance. The upper plots in this figure depict an example of the logistic decay with the $\alpha = 0.007$ and $d_0 = 500$m (right) and a grey scale heatmap plot (left) showing the aggregate capacity field in the plane with three WiFi base stations located as indicated by coloured cross markers. Similarly, the lower panel of Figure 4 shows an example of the logistic decay with the rate $\alpha = 0.07$ and displacement parameter $d_0 = 500$m, and the corresponding aggregate capacity field. Notably, the logistic functional form allows one to model *step-like* or *threshold* decays, as shown in the latter example (lower panel of Figure 4). We use the logistic decay, the locations of the WiFi base stations, and the *'clover'*-like trajectory from the upper left plot of Figure 4 in our further experiments in Section III.

In the next section we shall use the fluid approximation model from Section II-A1 for rate control in network topologies with time-varying channel capacities based on the above description. We use the iterative forward Euler method for solving the system of coupled differential equations (1) with time increment $\Delta t = 10^{-3}$s. (Note that $x_r(t)$ in (1) remains zero if the flow along route r stalls, even though there may be a chance for it to revive as channel capacities change. In such cases we imitate a *probing* strategy by setting $x_r(t-T_r)$ to a small positive constant of 10^{-1}s.)

III. Simulation Experiments

In this section we illustrate our fluid model to describe the behaviour of the primal form of the routing and congestion control algorithm discussed earlier in Section II.

The model parameters are given in Table I. The primal algorithm parameters were as follows: $\overline{a} = 0.025$, $b_r = 0.875$ for each router r. The resource responsiveness parameters β_j were chosen to be 15 and 16 for the six WiFi channels and satellite, respectively. (Note that this choice of parameters ensures local stability of the controller, as discussed in Section II-A3.) The nominal capacities C_j of the WiFi and satellite channels were all equal to $C = 20$ (e.g. 20 Mbps). Hererafter, we do not make any reference to a particular choice of physical units for nominal and actual capacities since all our results are scalable in terms of flow rates.

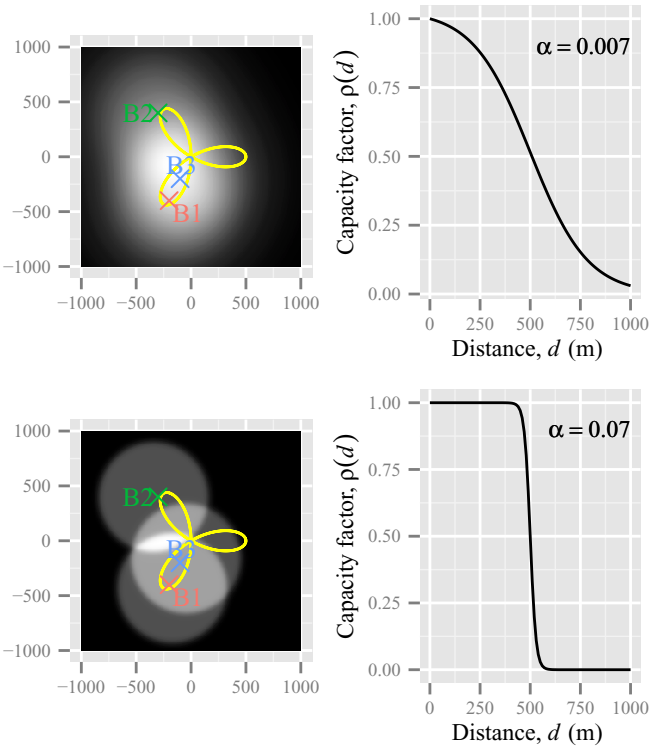

Fig. 4. **Top right:** Capacity factor as a function of distance using the logistic decay ($\alpha = 0.007$, $d_0 = 500$m). **Top left:** The aggregate capacity field in the plane with three WiFi base stations located as indicated by cross markers. We use the *'clover'* curve depicted in this figure to model a part of the trajectory of a moving convoy fleet in our experiments in Section III. **Bottom right:** Capacity factor as a function of distance using the logistic decay ($\alpha = 0.07$, $d_0 = 500$m). (Note the sharper *threshold*-like shape of the decay.) **Bottom left:** The aggregate capacity field in the plane with three WiFi base stations located as indicated by coloured cross markers.

TABLE I
Delay and round trip time parameters (in milliseconds)

j	r	T_{rj}	T_{jr}	T_r
j_1	r_1	15	25	40
j_2	r_2	10	10	20
j_3	r_3	10	10	20
j_4	r_4	15	15	30
j_5	r_5	15	25	40
j_6	r_6	10	10	20
j_7	r_7	10	10	20

A. Trajectory of motion and network resources

Consider the trajectory shown in Figure 5 of the central point of a convoy of length 200 metres (notice that the units of distance in Figures 4 and 5 are in correspondence and one unit of distance corresponds to one metre. Also shown are the fixed locations of three WiFi base stations labelled B1, B2 and B3. In our experiments the convoy, including its head, tail and middle point, follows the curve at a moderately variable speed ranging from approximately 8 m/s to 25 m/s (see Section II-B2 and Appendix for the precise details) passing the points $t_0, t_2, \ldots, t_{17}$ labelled by the instant the convoy passes that location. The time step from time t_i

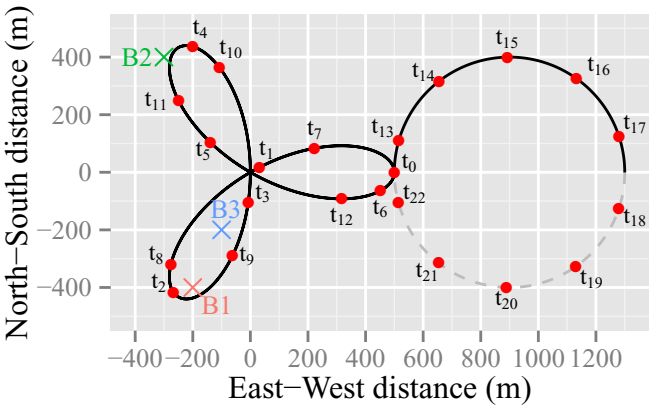

Fig. 5. This figure shows the trajectory of the moving convoy in a 2-D plane. The locations of the middle of the convoy at equal time steps of 30 seconds are labelled $t_1, t_2, \ldots, t_{17}$. The locations of three fixed base stations are shown by the coloured points labelled B1, B2 and B3 (blue, red and green, respectively). Note that the convoy makes two revolutions of the clover-shaped curve before moving away from the vicinity of the base stations along the circular arc towards the right hand side, and the entire trajectory (along the solid line) is completed in just under 9 minutes. The dashed lower arc of the circle is used in further experiments described in Section III-G.

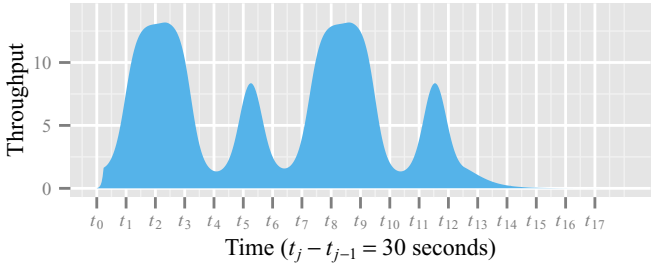

Fig. 6. This figure shows the throughput obtained by using a single WiFi channel to base station B1 as the convoy moves along the trajectory. Reference time epochs are labelled by $t_1, t_2, \ldots, t_{17}$ and shown by the vertical dashed lines. The throughput varies according to the relative location of the head of the convoy and B1.

to t_{i+1} represents 30 seconds. The convoy's head and tail positions move with respect to its middle point as discussed in Section II-B2 and shown in Figure 2. The convoy makes two complete revolutions of the clover shaped curve before then following the arc to the right hand side of the plane. As the convoy moves along the curve its distance to each base station will vary.

Additionally, a satellite channel may be available which connects the convoy to the remote destination via a channel of constant capacity (equivalently, the capacity factor for a satellite channel is taken to be $\rho(d) = 1$ for all $d \geq 0$).

B. Mobility without path diversity

In our first experiment we consider how the available throughput varied if the convoy had access via its head to just the single base station B1 while moving along the trajectory of Figure 5.

Figure 6 shows that the throughput is highly sensitive to the location of the convoy according to the time-varying distance between the head of the convoy and base station B1. Around times t_2 and t_8 the head of the convoy is closest to B1 and the throughputs adapt to take advantage of the available capacity at locations near to B1. In contrast, at around times t_4, t_6 and t_{10} the head of the convoy is distant from B1 and the flow is correspondingly low at or just after those epochs.

We can imagine that the head of the convoy could switch between base stations and thereby handover access via B1 to access via B2 or B3 according to whichever base station allowed the greatest throughput. Such a handover process would need to be explicitly coordinated and the handover itself would need to be finely judged to avoid the overhead of flipping between base stations offering very similar throughputs. In the next experiment we consider how the benefits of handover could be achieved without explicit coordination by means of MPTCP exploiting multiple paths with their diverse and time-varying properties.

C. Mobility with path diversity

In the second experiment we suppose that the head of the convoy can now use MPTCP to connect with the three base stations B1, B2 and B3. The primal algorithm will seek to allocate flow across the three routes and will update those flows according to the time-varying available channel capacities.

The upper panel of Figure 7 shows how the available capacity, $\rho(d)C_j$, varies over time for the channels to the three base stations. The bandwidth available to B2 is broadly out of sync with the bandwidth available to base stations B1 and B3 and the flows shown in the lower panel reflect this and allow the convoy to experience a smoother total flow except when it moves away from all three stations between times t_6 and t_7. This can also be seen in Figure 8 which shows the net flow together with individual flows via each base station. Here MPTCP allows the bandwidth to the three base stations to be aggregated. In the case of a hard handover strategy the head of the convoy would be connected to just one base station at a time.

Despite achieving a high degree of exploiting available bandwidth through connection along diverse paths the net bandwidth drops to zero once the convoy leaves the vicinity of the group of base stations after time t_{13}. In the next experiment we remedy this by use of the satellite channel that provides global coverage.

D. Mobility with additional path diversity

In the third experiment we have enhanced the convoy further by allowing the tail of the convoy to connect to the three base stations. The head and tail of the convoy are physically displaced being $L = 200$ metres apart (with respect to the path followed) and so can, with implicit coordination, take advantage of the available bandwidth to base stations afforded by different parts of the convoy.

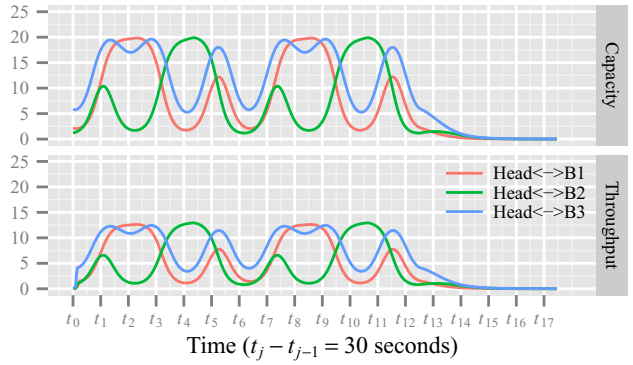

Fig. 7. This figure refers to the scenario in where the head of the convoy gains access through MPTCP to all three base stations (Section III-C). The upper panel shows how the capacity of each channel varies as the convoy moves along the trajectory and the lower panel shows the corresponding throughputs.

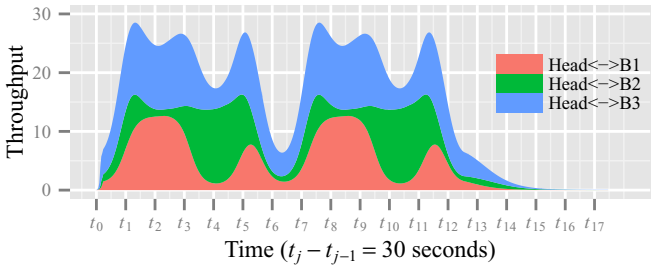

Fig. 8. This figure shows the individual throughputs over time for the experiment described in Section III-C. The shaded area refers to the throughput from the head of the convoy to B1 while the and areas refer to B2 and B3, respectively. Note that just after time point t_3 the throughputs to B1 and B3 are low since the convoy is a long distance from these base stations but the throughput to B2 has increased due to the closeness of the convoy to that base station. At around time point t_6 and from time point t_{13} onwards the convoy is not in the vicinity of any base station and the aggregate bandwidth suffers accordingly.

Figure 9 shows the capacities (upper panel) and throughputs (lower panel) over time for the six WiFi channels (three via the head of the convoy and three via the tail) and the single satellite channel, as in the diagram of Figure 3. While the convoy is in the vicinity of the base stations (up to time t_{12}) the satellite channel first reaches the level of around 12.8 and fluctuates slightly; in particular, we can see a slight rise just after time t_5 when there is a significant reduction in WiFi bandwidth. Beyond time t_{14} there is effectively no WiFi bandwidth available and the satellite channel is left to take over and is the sole supply of bandwidth.

The long-run throughput of the satellite channel is readily determined from the stationary solution, x_r^*, say, of the diffferential equations (1) where $\dot{x}_r(t) = 0$ for $r = 1, 2, \ldots, 7$. With no throughput on the WiFi channels in the long-run we have that $x_r^* = 0$ for $r \neq 4$ and the long-run satellite throughput, x_4^*, solves

$$\dot{x}_4(t) = 0 = \frac{x_4^*}{T_4}\,[\overline{a}(1-\lambda_4^*) - b_4 x_4^* \lambda_4^*]_{x_4^*}^{+} \qquad (12)$$

where we have used that $y_{s(4)}(t)$ becomes x_4^* in the long-run

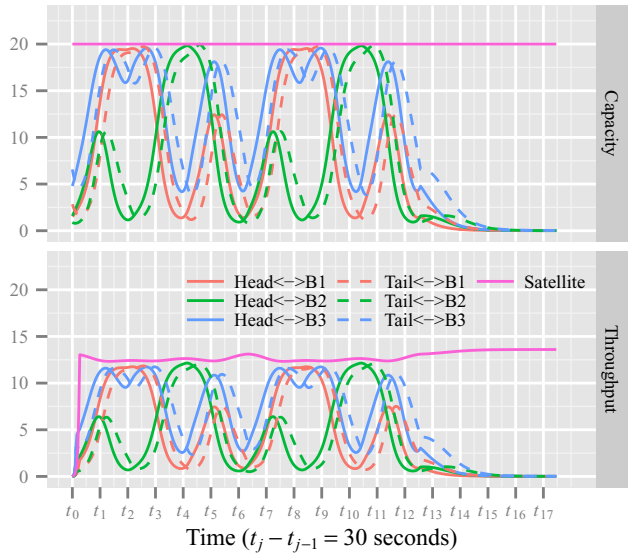

Fig. 9. This figure refers to the scenario where both the head and tail of the convoy can connect to each of the three base stations as well as to the satellite channel (hence, seven sub-flows). The upper panel shows how the channel capacities vary as the convoy moves along its route, and the lower panel shows corresponding channel throughputs. Note that the satellite channel can help support the convoy around time t_6 and after time t_{12} where the location of the convoy necessarily implies WiFi capacity is lacking.

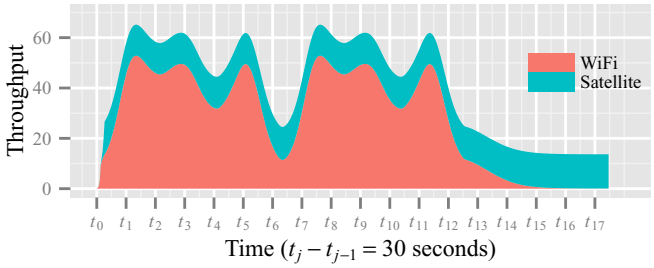

Fig. 10. This figure shows the WiFi bandwidth over time aggregated over the six channels as well as the satellite bandwidth . Due to the presence of the satellite channel there is always some available bandwidth and the coordination between the seven sub-flows has been achieved implicitly and smoothly without further explicit overheads.

and where by equation (5)

$$\lambda_4^* = p_{j_4}(x_4^*) = (x_4^*/C_4)^{\beta_4}\ . \qquad (13)$$

Thus we may solve for x_4^* from

$$\overline{a}(1-\lambda_4^*) - b_4 x_4^* \lambda_4^* = 0 \qquad (14)$$

and hence we obtain the fixed-point equation

$$x_4^* = \frac{\bar{a}}{b_4}\,\frac{1-\lambda_4^*}{\lambda_4^*}\,. \qquad (15)$$

This yields the equilibrium satellite channel throughput, $x_4^* = 13.6$ units (68% utilisation), which is in agreement with the numerical results of Figure 9. Some of our further experiments will also discuss how the utilisation of a channel could be improved (Sections III-F, III-G).

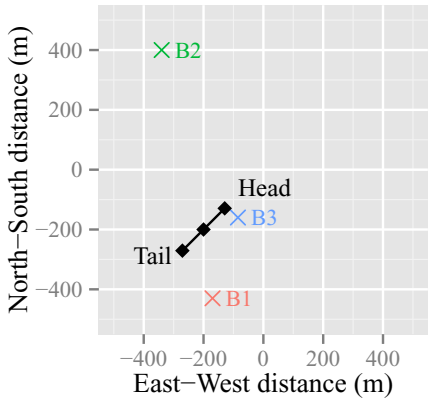

Fig. 11. This figure shows the location and orientation of the convoy as well as the WiFi base stations for the static experiment. Note that the WiFi base stations are at the same locations as in the moving scenario but the convoy is at a new location, not in fact on the trajectory of the moving convoy.

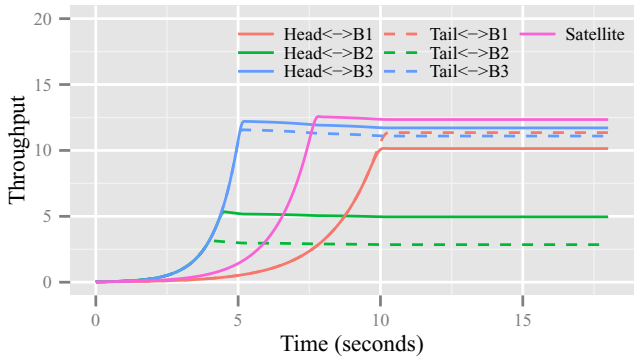

Fig. 12. This figure shows the individual sub-flows over time in the static scenario from Section III-E. There are seven sub-flows and they converge to their long-run levels at rates depending on the round trip times (see Table I).

E. Rates of convergence in static scenario

The experiments discussed so far have involved a moving convoy where the effects of the motion of the convoy and speed of convergence of the rate control algorithm are confounded. We now present a simple experiment where the convoy is static and investigate the convergence of just the rate control algorithm.

We suppose that the convoy is stationary at the location shown in Figure 11. Both the head and tail are close to base stations B1 and B3 and relatively far from base station B2. At this location the available capacities c_{j_k} are approximately given by $(16.97, 8.30, 19.59, 20, 18.99, 4.77, 18.56)$ units (see equations (9) and (10)).

Figure 12 shows the throughputs over time as the algorithm first increases and then levels off at the long-run throughputs within 5 to 10 seconds. Figure 13 shows how the total throughput is achieved according to the WiFi and satellite channels. Satellite is slower to converge than the WiFi associated with base station B3 but is faster to converge than the WiFi corresponding to communication via B1 due to the different round trip times (see Figure 3 and Table I).

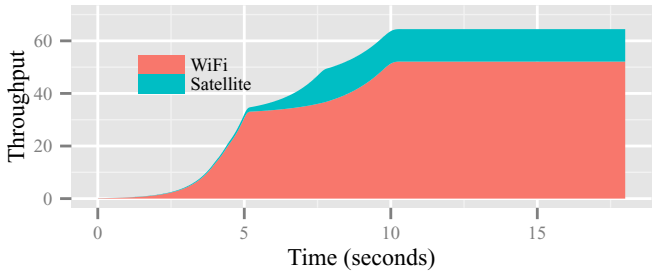

Fig. 13. This figure shows the equilibrium aggregated bandwidth achieved by the WiFi channels in comparison with the satellite bandwidth in the static scenario (see Section III-E).

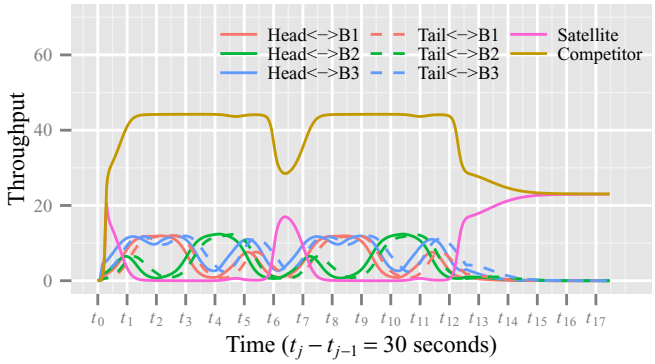

Fig. 14. This figure refers to the scenario in which there is a *competitor* for the satellite channel's bandwidth (see Section III-F). The figure shows all eight channel throughputs (seven fleet's ones and the *competitor*). Notably, and in contrast to Figure 9, the fleet's satellite channel supports the convoy more expressively when that is lacking WiFi access and is suppressed when the WiFi capacity is regained. The corresponding aggregated MPTCP throughputs are shown in Figure 15.

F. Sharing a channel

All experiments considered so far assumed that the communication channels used by the fleet were dedicated ones, that is the fleet did not have to share or compete for any of the available resources. However, it is not uncommon for transmission channels/resources to be shared; one such example is Multiplexed Channel Systems Demand Assignment Multiple Access (MCPC/DAMA) or "bandwidth on demand" systems [3].

In our next experiment we assume that the fleet shares the satellite channel with a "competitor". The experimental set-up is the same as in Sections III-B–III-E extended by adding a second communication source node s_2 to the system topology (Figure 3) which has access to the same satellite channel ($j = 4$). The system and controller's parameter values remain unchanged except for the capacity of this resource—because the controller implements the proportional fairness paradigm, we have scaled the satellite's channel capacity to $C_4 = 70$ units, a level at which the *competitor's* flow becomes comparable to the aggregate throughput for the fleet.

Figure 14 shows the throughputs over time in this experiment. As the satellite channel is shared by just two sources, the

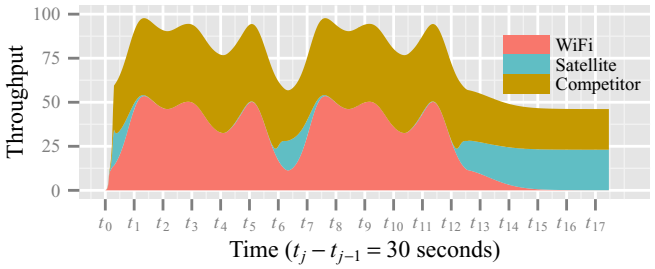

Fig. 15. This figure shows the throughput over time under the scenario when the satellite channel is shared with a *competitor* (Section III-F). The sub-flows of this communication system are represented as follows: WiFi aggregated flow , fleet's satellite channel , and a *competitor* for satellite .

fleet and *competitor*, their throughputs on this channel mirror each other. The fleet's satellite channel throughput appears much more responsive to changes in the aggregate WiFi flow than in our previous experiments — a feature which might be preferable from the fleet's point of view when, say, the satellite link is much more expensive to use than WiFi. This can also be seen in Figure 15 which shows the total WiFi, and fleet and *competitor* satellite throughputs. When there is no access to WiFi, the fleet and competitor share (a part of) the satellite resource equally and the fleet's throughput via this resource achieves its maximum (only slightly exceeding the dedicated satellite channel's capacity from previous experiments). This suggests that one way of achieving a better responsiveness of the throughputs on selected dedicated channels in multipath routing is by using *virtual* (higher than actual) channel capacities and additional (*virtual*) workload within the framework of MPTCP; of course, this can only be possible if the resource (e.g. a router) reveals to its end-hosts some internal network information, including the channel's maximally achievable throughput, which can be done, for example, using packet marking.

G. Utilisation and delay times

Our final experiment explores the effects of (i) improving utilisation by tuning controller's parameters, and (ii) higher round-trip times in environments with dynamically changing capacities.

In virtue of the discussion from Section III-D and equation (15) the channels' utilisation can be improved by increasing the responsiveness β_j of these resources (this decreases penalisation in the corresponding optimisation problem) and necessarily decreasing $\bar{a}$ in order to ensure stability. This, however, as well as high round-trip times, may contribute to a slower reaction to changing bandwidths.

The experiment is based on our main experimental set-up described in Section III-D. The trajectory which the convoy follows was extended to a full circular path after it leaves the clover curve having made two complete revolutions along it. All resource responsiveness values β_j were doubled, and $\bar{a}$ was halved; all one-way and round-trip times were taken tenfold higher than those given in Table I, and the ones which

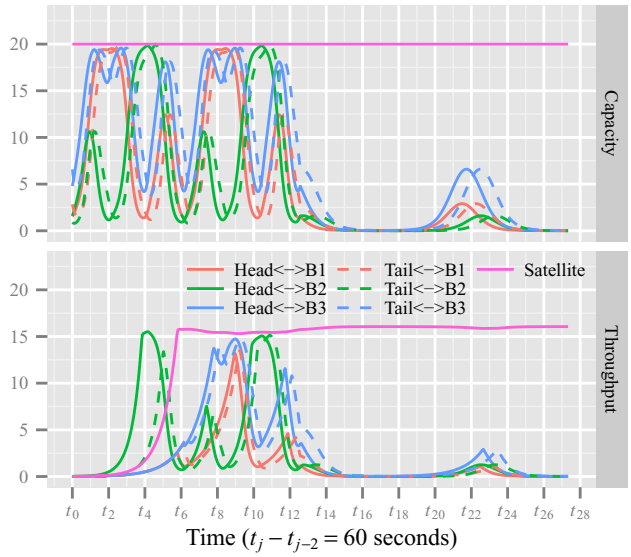

Fig. 16. This figure shows the WiFi channel capacities (upper panel) and all throughputs (lower panel) in an experiment described in Section III-G. This experiment is an extension of the one considered in Sections III-B-III-F with high resource responsiveness values and round-trip times. The trajectory of the fleet is also extended by a full circular path, as depicted in Figure 5 (both solid and dashed lines).

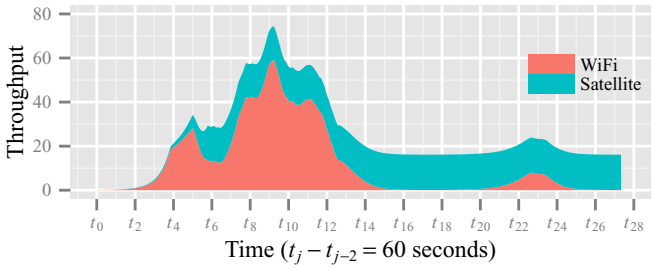

Fig. 17. This figure shows the bandwidth over time under fleet dynamic's scenario discussed in Section III-G. The aggregated flows in this figure correspond to the subflows from Figure 16 and are as follows: the WiFi aggregated flow of six channels and satellite channel .

correspond to the base station B_3 were further doubled.

Figure 16 shows the WiFi channel capacities and all seven throughputs. These feature much longer times to equilibrium compared to a similar experiment from Section III-D. The satellite throughput achieves its equilibrium level of 80.3% of the channel's bandwidth (see equation (15)) but hardly responds to drops (for example, between t_{12} and t_{18}) or increases (for example, between t_{12} and t_{23}) in aggregated WiFi throughput. Both the individual and aggregated WiFi throughputs (see Figure 17 for the latter) are now much less smooth than the corresponding channel capacities changing in time.

IV. Conclusions

This paper was motivated by the possibility of providing robust communication for a mobile convoy of vehicles through the exploitation of path diversity in a hybrid network. The

primal algorithm for a fluid approximation to proportionally-fair MPTCP model has been used to investigate the behaviour and patterns of sub-flows between a moving convoy and a remote destination via a hybrid network of WiFi base stations and a satellite channel.

The paper demonstrates how path diversity can be implicitly utilised to spread flow across available paths in both dedicated and shared channel systems. The pattern of observed sub-flows can be seen to be under dynamic control according to the motion of the convoy relative to the fixed base stations. The paper also discusses the issue of resource utilisation in such scenarios and highlights solutions for its improvement.

The work in this paper is influenced by theoretical discussions of rate control and routing algorithms operating at the fluid level and by recent demonstrations of the capabilities of MPTCP through packet level simulations. Our work is first in using fluid model approximation to explore some of the capabilities of MPTCP in dynamically changing settings and it describes the features and advantages of this approach that emerge implicitly without the overhead of explicit coordination. As the theory, practical implementations and standardisation of MPTCP controllers are still under development [8], [21], much further work remains to develop this approach as well as to capture the benefits in deployed protocols and study their behaviour under mobility scenarios.

ACKNOWLEDGMENT

This research was sponsored by the US Army Research Laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US Government, the UK Ministry of Defence or the UK Government. The US and UK Governments are authorised to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

The authors are thankful to Frank Kelly and colleagues within the Network and Information Sciences International Technology Alliance for useful discussions.

REFERENCES

[1] R. J. Gibbens, "Modelling multi-path problems," In *Proceedings of 42nd Annual Conference CISS 2008, Information Sciences and Systems*, pp. 42-45, 2008, Princeton, NJ, USA.

[2] H. Han, S. Shakkottai, C. Hollot, R. Srikant, and D. Towsley, "Overlay TCP for multi-path routing and congestion control," in *ENS-INRIA ARC-TCP Workshop*, 2003, Paris, France.

[3] "Handbook on Satellite Communications", *Int. Telecom. Union*, 3rd Ed., Wiley, 2002.

[4] F. Kelly and T. Voice, "Stability of end-to-end algorithms for joint routing and rate control," *Comput. Commun. Rev.*, vol. 35, no. 2, pp. 5–12, 2005.

[5] F. Kelly, "Fairness and stability of end-to-end congestion control," *Eur. J. Control*, vol. 9, pp. 159–176, 2003.

[6] T. Kelly, "Scalable TCP: improving performance in highspeed wide area networks," *Comput. Commun. Rev.*, vol. 32, no. 2, pp. 83–91, 2003.

[7] P. Key, L. Massoulié, D. Towsley, "Path selection and multipath congestion control," *Commun. ACM*, vol. 54, no. 1, pp. 109–116, 2011.

[8] R. Khalili, N. Gast, M. Popovic, U. Upadhyay, and J.-Y. Le Boudec. "MPTCP is not Pareto-optimal: Performance issues and a possible solution". In *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies*, Nice, France, 2012.

[9] Y. Lim, R. Gibbens, D. Towsley, "Multipath Transport in Hybrid Networks," In *'Hybrid Networks' Workshop of the Annual Conference of International Technology Alliance in Network and Information Sciences*, IBM US, September 27-28, 2011.

[10] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE ACM T. NETWORK.*, vol. 8, pp. 556–567, 2000.

[11] C. Paasch, G. Detal, F. Duchene, C. Raiciu, O. Bonaventure, "Exploring mobile/WiFi handover with multipath TCP,". In *Proceedings of the 2012 ACM SIGCOMM Workshop on Cellular Networks: Operations, Challenges, and Future Design*. ACM, New York, USA, pp. 31–36.

[12] F. Paganini, Z. Wang, J.C. Doyle and S.H. Low, "Congestion control for high performance, stability and fairness in general networks," *IEEE ACM T. Network.*, 2005.

[13] A. Papachristodoulou, L. Li, and J.C. Doyle, "Methodological frameworks for largescale network analysis and design," *Comput. Commun. Rev.*, vol. 34, no. 3, pp. 7–20, 2004.

[14] C. Raiciu, C. Paasch, S. Barre, A. Ford, F. Duchene, O. Bonaventure and M. Handley, "How hard can it be? Designing and implementing a deployable Multipath TCP," In *Procedings of Usenix NSDI*, 2012, San Jose, California.

[15] C. Raiciu, D. Niculescu, M. Bagnulo, M. Handley, "Opportunistic Mobility with Multipath TCP," In *Proceedings of the MobiArch Workshop (collocated with ACM Mobisys)*, 2011, Washington DC, USA.

[16] C. Raiciu, D. Wischik, M. Handley, "Practical congestion control for multipath transport protocols," *UCL Technical Report*, 2010.

[17] R. Srikant and S. Shakkottai, "Network optimization and control," *Found. Trends Netw.*, vol. 2, no. 3, pp. 271–379, 2007.

[18] R. Srikant, "The Mathematics of Internet Congestion Control," Birkhauser, 2004.

[19] D. Tse and P. Viswanath, "Fundamentals of Wireless Communication," Cambridge University Press, 2005.

[20] T., Voice, "Stability of Multi-Path Dual Congestion Control Algorithms," *IEEE ACM T. Network.*, vol. 15, no. 6, pp. 1231–1239, December 2007.

[21] D. Wischik, C. Raiciu, A. Greenhalgh, M. Handley, "Design, Implementation and Evaluation of Congestion Control for Multipath TCP," In *Proceedings of Usenix NSDI*, 2011, Boston, USA.

APPENDIX
FURTHER DETAILS ON CONVOY'S DYNAMICS

The parametrisation $\mathbf{S}(t) = (X(t), Y(t))$ of the "clover"-like trajectory used in our experiments is as follows:

$$X(t) = R\cos(3\gamma t)\cos(\gamma t)$$
$$Y(t) = R\cos(3\gamma t)\sin(\gamma t).$$

Since the length of the arc on this curve between points t_1 and t_2 is

$$\ell_{[t_1,t_2]} = R\gamma \int_{t_1}^{t_2} \sqrt{1 + 8\sin^2(3\gamma t)}\,\mathrm{d}t, \tag{16}$$

where the integrand ranges within the interval [1,3], then, measuring time in seconds and in order to ensure the maximum speed of 25 m/s, given the choice of $R = 500$ metres, the parameter γ was set to

$$\gamma = \frac{25\mathrm{m}}{3 \times 500\mathrm{m} \times 1\mathrm{s}} = 0.05/3\ \mathrm{s}^{-1}. \tag{17}$$

This choice of R and γ ensures the speed to range from approximately 8.5 m/s (such slow-downs occur at the far ends of the clover's petals) to 25 m/s (this maximal speed is achieved at the origin, see Figure 5) and gves a maximal acceleration of $0.84\mathrm{m/s}^2$. The convoy's speed is constant when it moves along the circle and is 8.5 m/s.