2009 American Control Conference
Hyatt Regency Riverfront, St. Louis, MO, USA
June 10-12, 2009

Global and Local Chassis Control based on Load Sensing

Mathieu Gerard, Michel Verhaegen

Abstract— **Modern cars are equipped with an increasing number of active systems in order to help the driver, e.g. to better deal with unexpected changes in the vehicle dynamics. To improve the global vehicle dynamics, comfort and handling, coordination between traditionally stand-alone active systems is required. This demands for a Global Chassis Controller.**

Thanks to the use of tyre force sensors, the controller proposed in this paper is structured on two layers which simplify the implementation and the computation.

The top layer deals with the global control and the force allocation. Thanks to the structure, the nonlinearities and uncertainties of the tyre do not appear in this layer and the allocation can be treated as a convex optimization problem. The solution to this problem taking tyre friction constraints into account is provided by a hybrid dynamical system.

The bottom layer features 4 independent local tyre controllers. These ensure, in a robust manner, that the allocated forces are realised.

A simulation example shows the performance of the method during a *split-mu braking* manoeuvre.

I. INTRODUCTION

Drivers are used to driving their cars in normal situations. Unfortunately, it can happen that the dynamics of the vehicle change drastically. Heavy braking, large yaw rate, low friction patches, large roll acceleration, tyre saturation, etc. are factors that make the vehicle exhibit unexpected reactions. Because the driver is not used to these reactions, he will have a hard time reacting properly and maintaining his or her desired trajectory. This leads to a decrease in comfort and safety. Many accidents take place because the driver ”loses control”.

One important way to make the vehicle safer and more comfortable or fun-to-drive [4] is to improve the predictability of the vehicle. This means that the car should also in extreme situations behave as in normal conditions, whatever its state or environmental disturbances. Even so, laws of physics can not be defied and therefore an option should be installed to warn the driver that the limit is approaching; this will be implemented in the future.

A. A virtual car

The objective of this research is to construct a control layer between the driver and the chassis that will make the driving more ”pleasant”: comfortable and safe [4]. Such a layer makes sure that the movements of the car follow precisely the instructions given by the driver. Unexpected car behaviours as well as disturbances will be directly compensated for before the driver can notice them. Finally the vehicle is controlled by two feedback loops. The inner-loop, closed by the controller, improves the dynamics of the mechanical system. The outer-loop, closed by the driver, gives the desired trajectory. The concept is illustrated on figure 1.

The inner-loop attempts to hide away to the driver the original car mechanics so that the driver controls a virtual car. The freedom to define the behaviour of such a virtual car depends upon the level of actuation. With configurations that are able to act on all the degrees of freedom, like the one used next, the possibilities are large. Using this approach, it becomes possible to decouple the design of the controller and of the virtual car. On one hand, experience on driver reactions can help defining the ideal virtual car, i.e. the car that drivers would love, based on comfort, safety and brand identity. On the other hand, control engineers ensure that the controlled car behaves like this ideal car.

B. Autonomous Corner Module

In order to have full freedom in controlling the car and defining the virtual car, only actuator configurations acting on all the degrees of freedom will be considered. Note that the degrees of freedom considered here are the longitudinal, lateral and yaw motions only. For the sake of brevity, only one configuration is discussed in this paper, but others are possible.

A tendency in vehicle design is to remove centralized actuators, like combustion engine and steering rack, and distribute them closer to the wheels. One example of such system is the so-called Autonomous Corner Module [11] where each wheel is fully actuated and electronically controlled. Each Corner Module is equipped with an electric motor acting as driving motor or as regenerative brake, with

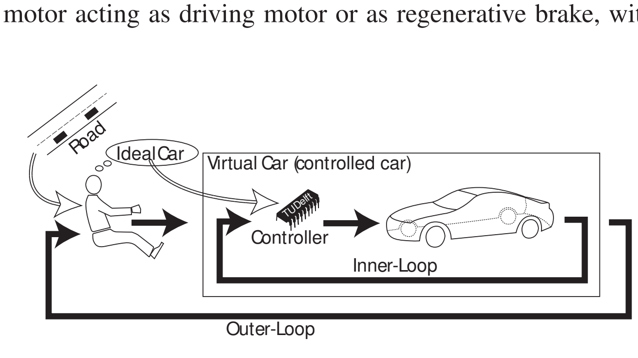

Fig. 1. The vehicle motion is controlled via two feedback loops. The inner-loop, closed by the controller, renders the car easier to drive. The tuning is based on the image the driver has about his ideal car. The outer-loop, closed by the driver, maintains the trajectory.

Mathieu Gerard and Michel Verhaegen are with the Delft Center for Systems and Control, Delft University of Technology. Mekelweg 2, 2628CD Delft, The Netherlands.

Corresponding Author: Mathieu Gerard m.p.gerard@tudelft.nl

a brake to provide larger braking torque than the regenerative brake and with a steering system. Other actuators like active suspension or active camber are not considered at this stage. On one hand, this configuration provides the most freedom for influencing the behaviour of the car. On the other hand, this is also the most challenging since it increases the complexity of the control law. A configuration with 4 corner modules will be discussed here.

C. Global Chassis Control

In this research, we assume that the driver instructions are already directly available in terms of desired longitudinal and lateral forces as well as a desired yaw moment for the chassis. The control problem is then to drive the actuators: wheel torque and steering, such that the generated forces and moment are the desired ones. To properly consider the coupling between the actuators, e.g. coming from the non-linear tyre dynamics, and to seek for global optimums for the control of the entire vehicle, a multi-input multi-output control strategy is necessary. In automotive, this has the name of Global Chassis Control.

The paper is structured as follows. In Section 2, the control problem will be presented in more detail and a structure is proposed. This structure will allow us to tackle the various difficulties with dedicated solutions instead of everything at once. Then, in Section 3, an original solution based on the hybrid steepest descent method will be applied to the force allocation part. A simple local tyre controller is introduced in Section 4. A simulation example is presented in Section 5. This simulation will show how the controller can help the driver brake on a road presenting a different friction coefficient on each side. The paper is concluded with some perspectives and future work.

II. CONTROLLER STRUCTURE

Let us recall that the objective of the controller is to drive the actuators in such a way that the forces generated at tyre level produces the desired total forces and moment on the chassis. A few major challenges can be identified:

- the system is over-actuated in the sense that there are more actuators than degrees of freedom,
- the system is non-linear and presents in particular saturations,
- many parameters, like the tyre-road friction at each tyre, are unknown and time varying.

A few solutions have been proposed in the literature. Borrelli [12] formulates a non-linear optimization problem in the Model Predictive Control framework. The actuator commands are directly computed by the algorithm. This can be seen as an advantage since actuator constraints can be easily formulated. However, the implementability is questionable because of the computational complexity of the method. Moreover, non-linear programming can lead to difficulties like local optima or slow convergence. Furthermore, many parameters and a tyre model have to be included in the model used in the optimization. Those are often very difficult to obtain in practice during real-time operation. Assessing the influence of inaccurate model parameters might be very complicated and including uncertainty in the non-linear MPC, will even be more computationally expensive.

To avoid solving non-linear programming problems online, Tøndel [15] proposes to solve the optimization offline using multiparametric programming. Solutions are then stored in large lookup-tables. This reduces the risk of bad convergence. However, the size of the lookup-tables can become impractical if the number of variables and parameters is large, which is the case in automotive applications. Furthermore, capabilities to reconfigure the system, for example in case of fault, are frozen.

Andreasson [13] tries another approach and solves at each time step the optimization problem based on a linearised model. This is motivated by the fact that the dynamics of the system should not vary too much between two iterations. Unfortunately, tyres can raise difficulties since they are known to have fast dynamics and strong non-linearities. Also, the computational complexity remains quite high.

Plumlee [14] simplifies the problem by linearising the vehicle model and therefore neglecting the non-linearities. Then a variant of Quadratic Programming is used.

The methods discussed above belong to the class of Direct Allocation since all the actuator commands are computed directly. Another possibility is to use Indirect Allocation where intermediate signals are used. These signals are chosen in order to simplify the allocation procedure, for example by making it convex, while allowing the actuator commands to be uniquely defined after some post-processing. As in [2], we propose to divide the controller in two layers: chassis and tyres. This is motivated by the facts that:

- chassis and tyre have different dynamics: the chassis (heavy mass) is slower than the wheel (light mass),
- chassis and tyre do not present the same non-linearities: the tyre has strong non-linearities that can make it go from stable to unstable,
- chassis and tyre do not have the same level of uncertainty: the inertia of the chassis is roughly known and does not vary during driving while the tyre-road parameters are time-varying and difficult to estimate.

In case Autonomous Corner Modules are used, the intermediate signals can be the longitudinal and lateral forces at each tyre, see figure 2. It can be noticed that, as long as the forces stay within the physical limits, they can be taken as independent: any combination can be achieved by appropriate steering angle and driving/braking force. By taking the intermediate forces in the vehicle frame (x direction pointing towards the front of the vehicle) instead of in the wheel frame (x direction steered with the wheel), the non-linearity linked to the steering of the wheels is placed in the second layer; rendering the optimization problem convex and far easier to solve.

III. DYNAMIC CONTROL ALLOCATION

The objective of this first layer is to distribute the total forces and moment desired on the chassis to the tyres, while respecting the tyre constraints. The total longitudinal and lateral forces and the yaw moment desired at the centre of gravity are given: $F^d = \begin{pmatrix} F_x^d & F_y^d & M_z^d \end{pmatrix}^T$. Let the forces at each of the 4 tyres be denoted by

$$F = \begin{pmatrix} f_{x_1} & f_{y_1} & f_{x_2} & f_{y_2} & f_{x_3} & f_{y_3} & f_{x_4} & f_{y_4} \end{pmatrix}^T \tag{1}$$

then the resulting forces and moment on the chassis are

$$\begin{pmatrix} F_x \\ F_y \\ M_z \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ -c & a & c & a & -c & -b & c & -b \end{pmatrix}}_{B} F \tag{2}$$

where a, b and c are constant parameters defining the position of the centre of gravity, see figure 2.

Allocating the forces then means that F should be computed such that $||BF - F^d||$ is as small as possible. In case constraints are active, it becomes necessary to weight the 3 terms in the minimization, i.e. to give priorities to the motions of the vehicle. For example, considering that the stability of the vehicle is the number-one priority, we could give a higher importance to the yaw moment. Furthermore, in case of emergency braking, we could assume that if the driver tries an evasive manoeuvre, this is for the safety of the vehicle, and it should therefore be fully assisted by giving higher priority to the lateral force compared to the longitudinal force. The way to tune these weights is currently under investigation. The weighting is done with the matrix W_R. However, the total forces and moments should be achieved by using the smallest possible individual tyre forces. This limits wear, keeps us away from the unstable region of the tyre, and it also avoids situations where forces are acting against each other. Again, a weighting between the 8 forces is required using the matrix W_F. For example, smaller weights can be given to the lateral forces compared to the longitudinal ones in order to corner using steering instead of differential braking.

Following Pacejka [7], the friction limit of the tyres takes the form of an ellipse. Since the parameters of that ellipse are rather uncertain and time varying, it is not necessary to take an accurate model, but an approximation is sufficient. In particular, it seems reasonable to approximate the ellipse by a rhombus [8]. During driving, the size of the rhombus is adjusted to match the locally identified tyre properties in the region of interest. The choice of a rhombus is preferred to a box in order to incorporate the combined-slip effect: at the friction limit, the force in one direction can only be increased if the force in the other direction is decreased.

This leads to the definition of the following optimization problem for F where the weighted 2-norm is used in the cost function and the subscript i in the rhombus constraints refers to the tyre number.

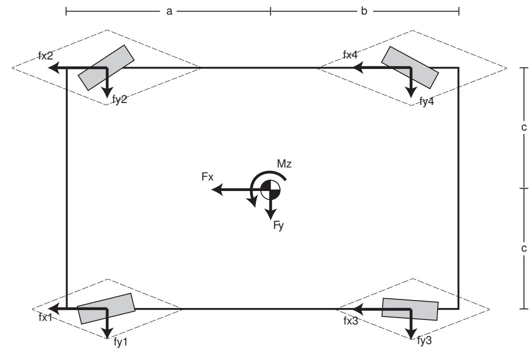

Fig. 2. Tyre force allocation from desired total forces and moment toward the tyres, within the vehicle frame.

$$\begin{aligned} \min_F \quad & \left\| \begin{matrix} F_x - F_x^d \\ F_y - F_y^d \\ M_z - M_z^d \end{matrix} \right\|_{W_R}^2 + \|F\|_{W_F}^2 \\ \text{subject to} \quad & -f_i^{\max} \le f_{x_i} + f_{y_i} \le f_i^{\max} \quad \forall i \\ & -f_i^{\max} \le f_{x_i} - f_{y_i} \le f_i^{\max} \quad \forall i \end{aligned} \tag{3}$$

Many methods have been presented in the literature to solve allocation problems. Good references can be found in [16] and [17]. Those methods are based on the idea that the precise solution to the allocation problem should be computed at each time step. However, this does not have to be the case. Since the speed of the actuators is limited, it is not needed to get the optimal value of F right away! What is important is that the value of F moves fast enough in the optimal direction, such that the actuators can follow, to finally end up at the optimal point. So we do not see F as a variable to be optimized over and over again, but as a variable that should be updated over time to minimize the cost function while staying in the acceptable region defined by the constraints. This can significantly reduce the computational complexity while a high level of performance is maintained.

Two methods have been developed very recently with this idea, see [1] and [9]. The methods differ in the way they handle the constraints. The method in [1] introduces the equality constraints in the cost function and uses the inequality constraints to define the feasible set, while [9] does the opposite: the inequality constraints are included in the cost function using barriers while the equality constraints define the acceptable region. Furthermore, in [9] the Lagrange multipliers are computed explicitly and need to be adapted on-line, while in [1] they are not computed but come as a by-product of the sliding mode. Finally, [1] is simple and easy to interpret. Both methods have stability and convergence proofs available. Therefore [1] is chosen for this application. Consequently, the update law of F takes the form

$$F(t + \Delta t) = F(t) + \sigma h(F(t), R, f_i^{\max}) \tag{4}$$

where σ is a tuning parameter defining the convergence rate

[3] and $h(.)$ is a discontinuous function of the form

$$h(F) = \begin{cases} -\nabla q(F) & \text{if } g_j(F) \leq 0 \quad \forall j \\ -\sum_{i \in L(F)} \nabla g_i(F) & \text{if } \exists j: \; g_j(F) > 0 \end{cases} \tag{5}$$

where the functions $q(F)$ and $g_i(F)$ are respectively the cost function and constraints in the optimization problem (3), and with $L(F) = \{l : g_l(F) \geq 0\}$.

The interpretation of the update law is the following [1]:

- as long as F is in the feasible set, the cost function is decreased using the gradient descent method;
- if F is outside the feasible set then F will be pushed back to the feasible set using a descent method for the constraint;
- if F is at the boundary of the feasible set, it might be pushed alternatively from one side to the other resulting in a sliding mode.

It is noted that the evaluation of $h(.)$ is always simpler than computing the value of the cost function and the constraints. Therefore, the practical implementation of such a method is easy and efficient.

IV. LOCAL TYRE CONTROLLER

The objective of each local tyre controller is to control one Corner Module via the steering system and the motor/brake such that the tyre develops the desired forces, i.e. determine the control inputs f_x and δ in order to drive the outputs F_x and F_y to the reference values, see figure 3. The approach should be as simple as possible and as robust as possible regarding the large uncertainty in tyre-road friction.

In the literature, some authors have proposed to use an inverse tyre model [2]. However, this solution is difficult to implement seeing the uncertainty on the tyre curve and the extreme difficulty to identify the full tyre model online. Moreover, this requires the computation of tyre slip, quantity which is still presently complicated to estimate accurately.

Our approach is based on the simple force feedback idea. Only the tyre forces need to be measured, for example using Force Sensing Bearings. In this paper, a very basic version of the controller is developed assuming that the tyre always stays in its stable region. This can be guaranteed for the longitudinal direction by using an ABS system. Such a simple controller is relevant to show the simplicity of this concept.

The tyre forces are measured in the bearing frame while the desired forces are expressed in the vehicle frame. Therefore, a correction is needed, see figure 3. For F_x and F_y the forces in the vehicle frame, f_x and f_y the forces in the bearing frame and δ the steering angle, we have the following equations:

$$\begin{pmatrix} F_x \\ F_y \end{pmatrix} = R \begin{pmatrix} f_x \\ f_y \end{pmatrix} \text{with} R = \begin{pmatrix} \cos(\delta) & -\sin(\delta) \\ \sin(\delta) & \cos(\delta) \end{pmatrix} \tag{6}$$

Our first actuator is the motor/brake. The longitudinal dynamic of the tyre being very fast, we can assume that the transfer function between the driving/braking torque and the driving/braking force is simply a static gain approximately equal to $\frac{1}{r}$ for r the radius of the wheel. Therefore the applied torque is taken r times the desired longitudinal force.

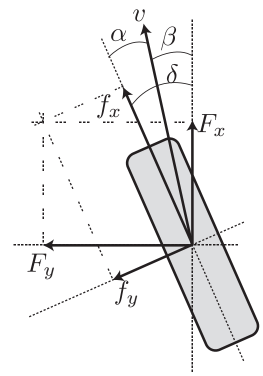

Fig. 3. Tyre forces are measured in the bearing frame while the desired forces are expressed in the vehicle frame.

The second actuator is the steering system. We take a static tyre model in differential form where only the main dependency in α is taken into account:

$$\begin{aligned} \dot{f}_y &= C_\alpha \dot{\alpha} && (7) \\ &= C_\alpha(-\dot{\beta} + \dot{\delta}) && (8) \end{aligned}$$

where $C_\alpha = \frac{\partial f_y}{\partial \alpha}$ is the local cornering stiffness which can vary depending the operating point, α is the lateral slip angle, β is the body slip angle and δ is the steering angle.

The relation between F_x and F_y, and f_x and δ in differential form can be obtained using the chain rule and equations (6) and (8)

$$\begin{pmatrix} \dot{F}_x \\ \dot{F}_y \end{pmatrix} = \begin{pmatrix} \frac{\partial F_x}{\partial f_x} & \frac{\partial F_x}{\partial \delta} \\ \frac{\partial F_y}{\partial f_x} & \frac{\partial F_y}{\partial \delta} \end{pmatrix} \begin{pmatrix} \dot{f}_x \\ \dot{\delta} \end{pmatrix} \tag{9}$$

with

$$\begin{aligned} \frac{\partial F_x}{\partial f_x} &= \cos(\delta) && (10) \\ \frac{\partial F_x}{\partial \delta} &= -f_x \sin(\delta) - C_\alpha \sin(\delta) - f_y \cos(\delta) && (11) \\ \frac{\partial F_y}{\partial f_x} &= \sin(\delta) && (12) \\ \frac{\partial F_y}{\partial \delta} &= f_x \cos(\delta) + C_\alpha \cos(\delta) - f_y \cos(\delta) && (13) \end{aligned}$$

Since β is a priori not known, it is regarded as a disturbance and removed from the nominal model.

A simple integral controller can be constructed:

$$\begin{pmatrix} \dot{f}_x \\ \dot{\delta} \end{pmatrix} = \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix} \left[R^T \begin{pmatrix} F_x^d \\ F_y^d \end{pmatrix} - \begin{pmatrix} f_x \\ f_y \end{pmatrix} \right] \tag{14}$$

Using this controller, the closed-loop system can be rewritten in state-space form:

$$\begin{aligned} \begin{pmatrix} \dot{F}_x \\ \dot{F}_y \end{pmatrix} &= \begin{pmatrix} \frac{\partial F_x}{\partial f_x} & \frac{\partial F_x}{\partial \delta} \\ \frac{\partial F_y}{\partial f_x} & \frac{\partial F_y}{\partial \delta} \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix} R^T \begin{pmatrix} F_x^d - F_x \\ F_y^d - F_y \end{pmatrix} \\ &= \underbrace{R \begin{pmatrix} a & -b f_y \\ 0 & b(f_x + C_\alpha) \end{pmatrix}}_{A} R^T \begin{pmatrix} F_x^d - F_x \\ F_y^d - F_y \end{pmatrix} && (15) \end{aligned}$$

This is clearly a linear time-varying system. Using Lyapunov arguments, the stability of this system can be assessed by looking at the eigenvalues of $A + A^T$ [18]. Acting as a congruent transformation, R does not have any influence on the stability. In our case, it is straightforward to check that the system is robustly stable as long as we stay away from the limits of the tyre.

As can be seen from its structure, the controller contains an integral action for both forces. Therefore the tracking error will be reduced to zero whatever the disturbance or the uncertainty in the system. The method tried on nonlinear tyres including roll-resistance and relaxation effects gives very good results, as can be seen in next section. Moreover, no tyre parameter needs to be known or estimated. If a rough estimate of the body slip angle β is known, it can be used as feedforward on δ to speed-up the system, but this is not necessary.

Of course, such a simple controller has room for improvement in later versions. In particular, future work will focus on robust controller design taking into account all the dynamics of the tyre like relaxation, actuator dynamics like brakes or steering systems and tires limits.

V. MODELLING AND SIMULATION

A 3D full car model is developed in Dymola based on the VehicleDynamics Library [5]. Dymola is an object-oriented physical modelling environment. It allows the simulation of any equation-based model and in particular of multi-body models. The VehicleDynamics Library exploits this multi-body capability to construct a complete car model with a high level of detail.

Our vehicle uses the *body* object, the equations of motion, the tyre models and the road contact calculation from the library. Then the traditional suspensions are replaced by four corner modules and the global and local controllers are added.

The simulation is based on the scenario of *split-mu braking*, which means that the driver wants to brake in straight line while the left side of the road is on a more slippery surface (like ice or snow) than the right side (asphalt). Without control, it is easy to understand that the brake efficiency will be larger on the right than on the left and therefore the car will start turning to the right. To keep the trajectory, the driver would need to compensate for that, which might be difficult and unsafe.

Our objective is to show that using a global chassis controller, the driver will be able to perform the manoeuvre without noticing that the situation is not ”normal”, which improves the drivability of the car, the comfort and in particular the safety.

The reference given by the driver is the following: $F_x^d = -3000$ N, $F_y^d = 0$ and $W_z^d = 0$.

The desired longitudinal and lateral forces are directly used in the allocation method while the desired yaw moment M_z comes from a PI controller maintaining the yaw rate W_z at the desired value.

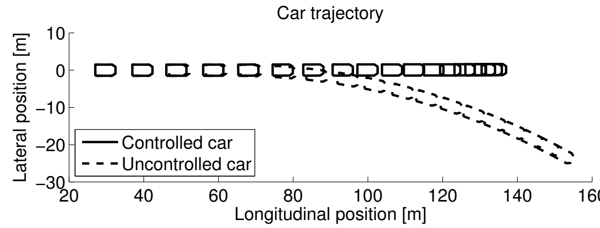

Fig. 4. Comparison with the uncontrolled case. Without control, one forth of the total braking torque is applied to each wheel. Note that because of the limited friction, only 100N is generated on the left. Without control, an undesired lateral displacement is clearly visible.

Because of the tyre-road friction characteristic, the maximum force reachable on a left tyre is 100 N. In this simplified implementation, the constraint is fixed before-hand. Future work includes the detection of this limit on-line. This can be done by detecting the moment when the tyre cannot reach the reference any more.

The results of the simulation are presented in figure 4 and figure 5. Figure 4 shows the difference between the controlled and the uncontrolled case. In the uncontrolled case, an undesired lateral displacement is clearly visible. Figure 5 focusses on the controlled case. The first plot displays the longitudinal and lateral acceleration as well as the yaw rate. The second plot presents the forces measured at the tyres in the longitudinal direction of the vehicle (not the longitudinal direction of the tyre). The last plot shows the forces measured at the tyres in the lateral direction of the vehicle.

The first 2 seconds are reserved for the initialization of the multi-body model. It can be noticed that the longitudinal acceleration is not zero because of the rolling resistance. At time $t = 2$ seconds, the controllers are turned on and the local controllers start compensating for the rolling resistance in order to get a zero longitudinal acceleration. The bias in the measured longitudinal forces is due to the misalignment of the force sensors because of the non-zero pitch angle (the car is heavier at the front). Moreover, the local controllers will slightly steer the wheels to compensate for the asymmetry of the tyre and the toe-in. At time $t = 3$ seconds the driver applies the brake command and the reference for the longitudinal force becomes $Fx = -3000$ N. Instantaneously, the wheels start braking but the left wheels are limited at 100 N because of the low friction. To compensate, the global controller asks for lateral forces on the right tyres. The left tyres cannot provide any lateral force because of the combined-slip effect. Therefore, the front-right wheel will turn to the left while the rear-right wheel will steer to the right. After a small transient, both the lateral acceleration and the yaw rates come back to zero while the braking command is perfectly satisfied.

VI. CONCLUSION AND FUTURE WORK

In this paper, a framework for Global Chassis Control of a fully actuated car has been presented. Thanks to the two layer structure and the intermediate distribution of tyre forces in the vehicle frame, the allocation problem is rendered convex,

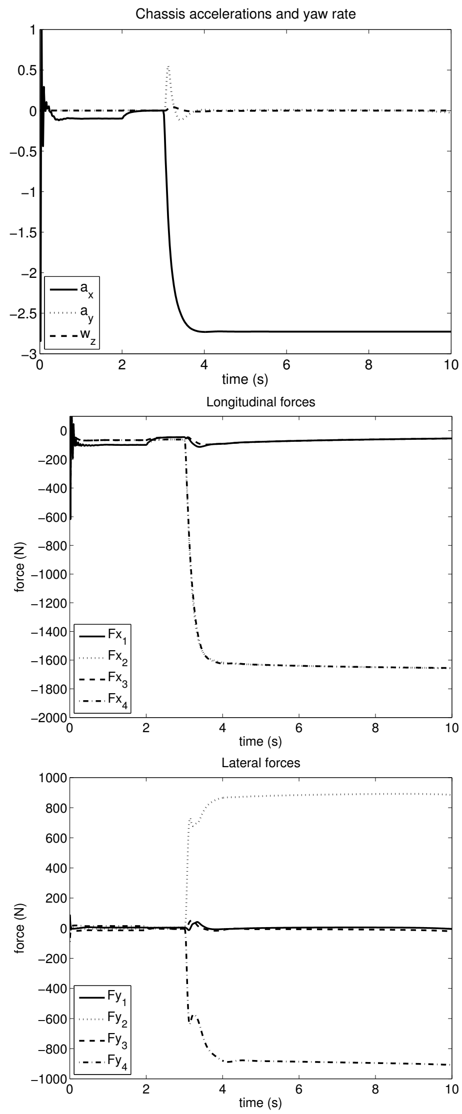

Fig. 5. Result of the simulation for a controlled split-mu braking manoeuvre. Because of low friction, the forces on the left tyres are limited to 100 N while the asked total braking force is -3000 N. Thanks to the redistribution of forces, the driver does not notice anything from the dangerous situation.

i.e. easier to solve, with no nonlinearities and no parameter uncertainties. The problem of selecting the weighting parameters in the allocation cost function still remain to be tackled. It is, for example, difficult to know what priority to give to certain movements in critical situation, when not all movements can be achieved.

Then, a new dynamic allocation method based on the Hybrid Steepest Descent Method [1] has been applied. This method gives excellent first results in the simulation study and has the advantage of being very easy to implement and to interpret.

Further, a very simple integral controller is designed at tyre level. It can be noticed that finally, not a single tyre parameter is required to perform accurate control. Tyre slip is also not used and therefore an accurate estimation of the longitudinal vehicle speed is not necessary. Still the local tyre controllers will be improved in future work.

Obviously, the whole architecture is based on tyre force measurement, but we have good reasons to believe that such sensors will arrive soon on the market.

The method is shown to perform very well during a *split-mu braking* manoeuvre. If a friction limit is detected on a tyre, the allocation method will automatically redistribute the forces in order to maintain the desired motion. Therefore, the unsafe environment becomes completely invisible to the driver, at least before approaching the physical limit. This improves both comfort and safety.

VII. ACKNOWLEDGMENTS

The authors would like to thank the Competence Centre for Automotive Research (CCAR) for funding this research.

The frequent discussion with experts from the automotive research labs of SKF (Bernie van Leeuwen) and TNO (Jeroen Ploeg and Jan Pieter Maurice) has been greatly appreciated.

REFERENCES

[1] M. Gerard, M. Verhaegen, B. De Schutter. A Hybrid Steepest Descent Method for Constrained Convex Optimization. Submitted for publication. 2007.

[2] J. Andreasson. On Generic Road Vehicle Motion Modelling and Control. Doctoral Thesis in Vehicle Dynamics. Stockholm, Sweden. 2006.

[3] M. Gerard, M. Verhaegen. Model Predictive Control using Hybrid Feedback. IFAC World Congress, Seoul, Korea. 2008.

[4] C. Canudas-de-Wit, H. Bechart, X. Claeys, P. Dolcini, J-J. Martinez. Fun-to-Drive by Feedback. International Journal of Control, special edition ECC/CDC 05. 2005.

[5] J. Andreasson, M. Gäfvert. The VehicleDynamics Library — Overview and Applications. International Modelica Conference, Vienna, Austria. 2006.

[6] L. Laine. Reconfigurable Motion Control Systems for Over-Actuated Road Vehicles. Doctoral Thesis, Department of Applied Mechanics, Chalmers University of Technology. Gteborg, Sweden. 2007.

[7] H. Pacejka. Tyre and Vehicle Dynamics. Elsevier, Oxford, UK. 2006.

[8] B. Schofield, T. Hägglund, A. Rantzer. Vehicle Dynamics Control and Controller Allocation for Rollover Prevention. IEEE International Conference on Control Applications. 2006.

[9] T. Johansen. Optimizing Nonlinear Control Allocation. IEEE Conference on Decision and Control, Bahamas. 2004.

[10] J. Tjønnås, T. Johansen. Adaptive Optimizing Dynamic Control Allocation Algorithm for Yaw Stabilization of an Automotive Vehicle using Brakes. 14th Mediterranean Conference on Control and Automation. 2006.

[11] M. Jonasson, S. Zetterström, A. Trigell. Autonomous Corner Modules as an Enabler for New Vehicle Chassis Solution. FISITA World Automotive Conference, Yokohama, Japan. 2006.

[12] F. Borrelli, P. Falcone, T. Keviczky, J. Asgari, D. Hrovat. MPC-based approach to active steering for autonomous vehicle systems. International Journal of Vehicle Autonomous Systems. Vol 3, Nb 2-4, Pg 265 - 291. 2005.

[13] J. Andreasson, T. Bünte. Global Chassis Control based on Inverse Vehicle Dynamics Models. Journal of Vehicles Systems Dynamics, vol. 44. 2006.

[14] J. Plumlee, D. Bevly, A. Hodel. Control of a Ground Vehicle using Quadratic Programming Based Control Allocation Techniques. American Control Conference, Boston, USA. 2004.

[15] P. Tøndel, T. Johansen. Control Allocation for Yaw Stabilization in Automotive Vehicles using Multiparametric nonlinear programming. American Control Conference, Portland, USA. 2005.

[16] M. Bodson. Evaluation of Optimization Methods for Control Allocation. Journal of Guidance, Control, and Dynamics, vol. 25, no. 4, pp. 703-711, July 2002.

[17] O. Härkegård. Backstepping and Control Allocation with Applications to Flight Control. PhD thesis, Department of Electrical Engineering, Linköping University, Sweden. 2003.

[18] Slotine, J.J.E., and Li, W., Applied Nonlinear Control, Prentice-Hall, 1991.