Copyright © IFAC Large Scale Systems,
Rio Patras, Greece, 1998

LARGE SCALE SYSTEM ISSUES IN AUTOMATED HIGHWAY AND AIR TRAFFIC MANAGEMENT SYSTEMS *

John Lygeros, George J. Pappas and Shankar Sastry

Department of Electrical Engineering and Computer Sciences
University of California, Berkeley, CA 94720
lygeros,gpappas,sastry@eecs.berkeley.edu

Abstract: Large scale systems typically consist of compositions of many interacting sub-systems performing various planning and control functions. Control for such systems is often arranged in a hierarchical structure, with high level *scheduling functions* concerned with optimality and efficiency issues and low level *control functions* implementing the desired high level objectives and dealing with system uncertainty and safety. This paradigm leads to interactions between the fields of optimization and control theory. In this paper, these issues are highlighted using two large scale applications: Automated Highway Systems (AHS) and Air Traffic Management (ATM) Systems. *Copyright © 1998 IFAC*

Keywords: Large-scale systems, hierarchical control, traffic control, air traffic control, safety critical systems.

1. INTRODUCTION

The term Large Scale Systems (LSS) is often used to refer to systems with multiple components that need to cooperate to achieve a common goal. Typical examples include process control systems, manufacturing systems, power generation and distribution systems and transportation systems (such as highway and air traffic management systems). The last two examples are sometimes also referred to as *multi-agent systems*, as they consist of a large number of agents (vehicles or aircraft) trying to make efficient use of a common resource (highways, airports, and airspace). Most of the systems discussed above are safety critical; small flaws in their design may result in substantial risk to human lives and/or the environment. Moreover, the size of the systems is such that even minor improvement in their performance may lead to substantial financial benefit. There is therefore a clear need to increase both the *safety* and *efficiency* of LSS.

The complexity of LSS often forces designers to break up the controller into a *hierarchy*. In most cases the functionality of the controller can be macroscopically decomposed as the two level hierarchy shown in Figure 1. The lower level includes controllers whose primary objective is to maintain safe operation; the techniques used to design this level typically come from the linear, nonlinear and hybrid control. The high level controller is responsible for scheduling and optimization and its primary objective is to improve the efficiency of the system. The tools used in this level come from operations research, mathematical programming and queuing theory.

In many LSS the requirements for safety and efficiency may be in conflict. This is especially pronounced in multi-agent systems; typically safety

* Work supported by NASA grant NAS 2-14291, ARO grants DAAH 04-95-1-0588 and DAAH 04-96-1-0341, PATH grants MOU 238, and MOU 319, and ONR grant N0014-97-1-0946.

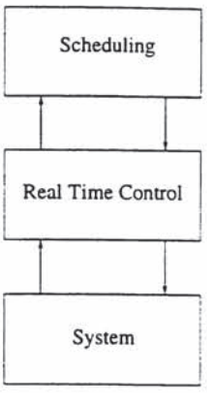


Fig. 1. Functional Hierarchical Decomposition

requires keeping the agents physically separated whereas efficiency requires brining them closer together. The levels of the hierarchy have to achieve an acceptable trade off between these two conflicting requirements. Our approach to this problem has been to try to establish bounds under which the low level controller can meet the safety requirements, using optimal control and game theoretic tools. Then, at the high level, we seek to optimize the efficiency of the system within these bounds. Note that this approach implicitly assumes that safety is given higher priority than efficiency.

The interplay between safety and efficiency considerations is too general to address in a theoretical setting. Here we use two applications to motivate the relevant issues. The first application, presented in Section 2, is Automated Highway Systems (AHS). Game theory is used to quantify the limits of safe operation. Linear programming techniques can then be used to maximize the highway throughput within these limits. The second example is Air Traffic Management (ATM) systems (Section 3). Motivated by the Center-TRACON Automation System (CTAS) and Flight Management Systems (FMS), we discuss how optimization and control considerations can come together to generate safe schedules and provide advisories that allow air traffic controllers and pilots to implement them.

2. AUTOMATED HIGHWAY SYSTEMS

A proposed solution to improve the efficiency of current highways is the the concept of platooning (Varaiya, 1993), where vehicles travel in tightly spaced groups (platoons), following each other with spacings of the order of 1 meter. Platoons are separated by large spacings, of the order of 60 meters. A control hierarchy for organizing traffic in platoons was proposed by Varaiya in (Varaiya, 1993). The hierarchy consists of 4 layers. The top two form the *roadside controller* and play the role of the "scheduling" level of Figure 1. Their

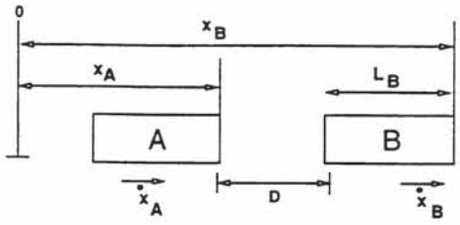


Fig. 2. Vehicle Following

goal is to maximize the throughput of the highway system by monitoring the aggregate flow of traffic and broadcasting commands to all the vehicles. The bottom two layers of the hierarchy form the *on board controller* and are responsible for "real time control". Their primary goal is safety. To achieve this goal they monitor the position of a vehicle with respect to its neighbors and issue commands to its actuators.

In recent years controllers have been designed for all the layers of the hierarchy. Here we outline a methodology (Lygeros *et al.*, 1998*a*) for establishing conditions under which the operation of the system can be safe and for designing on board controllers that guarantee safety. We then show how the safety requirements produced by this methodology can lead to estimates of the throughput of the automated highway.

2.1 *Real Time Control*

Consider two vehicles (labeled A and B) moving along a single lane highway (Figure 2). Assume that vehicles A and B have lengths L_A and L_B and let x_A and x_B denote their positions with respect to a fixed reference on the road. We seek a control scheme for vehicle A and requirements on the initial conditions and the system parameters such that no collision is possible between the two vehicles, whatever vehicle B decides to do. Following (Lygeros *et al.*, 1998*a*), we treat this problem as a game between the actions of vehicle A and those of vehicle B.

Assume that vehicle A can be modeled by a third order system and that the acceleration of vehicle B, denoted by d, can not be measured by vehicle A. If we let $D = x_B - x_A - L_B$ the system A-B can be described by the state vector $x = [\dot{x}_A \ \ddot{x}_A \ D \ \dot{D}]^T$. After feedback linearization (Godbole and Lygeros, 1994) the evolution of the state is described by:

$$\dot{x} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & -1 & 0 & 0 \end{bmatrix} x + \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} u + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} d \quad (1)$$

where u is the jerk applied by the controller of vehicle A. Let x^0 denote the value of the state at

time $t = 0$. The vehicle dynamics are constrained by the engine, tire and road conditions. More specifically it is required that $x_1 \in [v^A_{min}, v^A_{max}]$, $x_2 \in [a^A_{min}, a^A_{max}]$, $x_4 + x_1 \in [v^B_{min}, v^B_{max}]$, $u(t) \in [j_{min}, j_{max}]$, and $d(t) \in [a^B_{min}, a^B_{max}]$. This model is adequate for the case where there is no communication of state information between the platoons and no sensing and actuation delays. For an extension to the analysis presented here to different models see (Godbole and Lygeros, 1997).

For the purposes of safety we would like vehicles to avoid collisions. For vehicle A, this requirement can be encoded by a cost function:

$$J(x^0, u, d) = -\inf_{t \geq 0} x_3(t) \tag{2}$$

If, for a given initial condition x^0 and a given choice of u and d, $J(x^0, u, d) \leq 0$ m, vehicle A will never collide with vehicle B. We would like vehicle A to remain safe in this sense whatever vehicle B decides to do. To establish conditions for which this is possible we seek the worst possible action of vehicle B and the best possible action of vehicle A. In other words we seek a pair (u^*, d^*) such that:

$$J(x^0, u^*, d) \leq J(x^0, u^*, d^*) \leq J(x^0, u, d^*)$$

Such a choice of inputs is called a *saddle strategy* to the two player, zero sum game between u and d with cost J. Consider the candidate feedback saddle strategy:

$$u^* = \begin{cases} j_{min} & \text{if } x_2 > a^A_{min} \\ 0 & \text{otherwise} \end{cases} \tag{3}$$
$$d^* = \begin{cases} a^B_{min} & \text{if } x_1 + x_4 > v^B_{min} \\ 0 & \text{otherwise} \end{cases}$$

The candidate saddle simply requires both vehicles to brake as hard as possible. Let x^* denote the state trajectory under input (u^*, d^*) and $J^*(x^0) = J(x^0, u^*, d^*)$. In (Lygeros *et al.*, 1998*a*) it was shown that:

Lemma 1. (u^*, d^*) is globally a saddle strategy for cost $J(x^0, u, d)$.

Let T_3 be the time when $x^*_1(T_3) = 0$ (vehicle A stops). Note that $x^*_3(t)$ is a continuously differentiable function of t defined on the compact interval $[0, T_3]$, with derivative $x^*_4(t)$. Therefore:

Lemma 2. There exists $\hat{T} \in [0, T_3]$ such that $J^*(x^0) = -x^*_3(\hat{T})$ and either $\hat{T} = 0$ or $\hat{T} = T_3$ or $x^*_4(\hat{T}) = 0$.

Lemma 2 allows us to calculate $J^*(x^0)$ for a given initial condition x^0. In particular it allows us to distinguish safe situations ($J^*(x^0) < 0$) from unsafe ones ($J^*(x^0) > 0$) and determine the boundary between them ($J^*(x^0) = 0$). The boundary can be computed by a function, s, that returns the minimum inter-vehicle spacing needed for safety for given values of the velocity, acceleration and relative velocity of the vehicles:

$$s : \Re^3 \longrightarrow \Re$$
$$(x^0_1, x^0_2, x^0_4) \longmapsto x^0_3$$

such that $J^*([x^0_1 \; x^0_2 \; s(x^0_1, x^0_2, x^0_4) \; x^0_4]^T) = 0$. The boundary provides a sufficient condition for a feedback controller to be safe; any controller that applies u^* whenever $J^*(x(t)) = 0$ is guaranteed to be safe, as long as $J^*(x^0) \geq 0$.

The above methodology can be applied to all the controllers required for platooning operation: for the lead vehicle of a platoon, for the followers, for joining and splitting platoons, for lane changing, for entering and for exiting. In (Lygeros *et al.*, 1998*a*) it was shown such controllers can be designed and how the discrete switching between them can be implemented in a way that guarantees the safety of the overall, hybrid system.

2.2 *Throughput & Scheduling*

The minimum inter-vehicle safe spacing gives rise to what is known in the transportation literature as the *pipeline throughput.* Pipeline throughput is defined as the steady state per lane capacity of the highway without considering flow disturbances due to lane changes, entry and exit. Assume an AHS is populated entirely by platoons of size N moving at velocity x^0_1. Let F denote the follower spacing and assume all vehicles have length L. The throughput that can be achieved in safety on such an AHS (in vehicles per lane per unit time) is given by:

$$Q = \frac{N x^0_1}{s(x^0_1, 0, 0) + NL + (N-1)F} \tag{4}$$

This is clearly a very optimistic estimate. The actual value of the throughput is likely to be a lot smaller because of safety margins and disturbances introduced by entering and exiting vehicles, lane changes, etc. Even though it is optimistic, the above estimate can be treated as an upper bound on the throughput that can be achieved. For example, assume that the vehicles have no knowledge of the deceleration capability, a^A_{min}, but know that it is bounded in an interval $[\underline{a}, \overline{a}]$. We can model the fact that the deceleration of the leader is limited by that of the followers by assuming that platoons follow each other at spacings of:

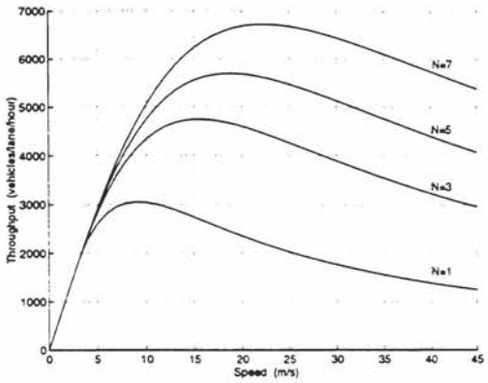


Fig. 3. Pipeline throughput for platoons of size N

$$s_{max} = \max_{\gamma a^A_{min}, a^B_{min} \in [\underline{a}, \overline{a}]} s_L(x_1^0, 0, 0)$$

where $\gamma = 1, 1.05, 1.1, 1.15, 1.2$ for $N = 1, 2, 3, 4$ and $N \geq 5$ respectively [1]. The resulting throughput for $\underline{a} = -9.31ms^{-2}$, $\overline{a} = -4.95ms^{-2}$ $F = 1m$ and $L = 5m$ is shown in Figure 3.

The estimate can be refined to allow for sensing and actuation delays, safety margins and different information structures (Godbole and Lygeros, 1997). The results can also take into account entry/exit disturbances, lane changes, etc. In (Broucke and Varaiya, 1996) a roadside controller was proposed to optimize the throughput of the highway by appropriately scheduling vehicle maneuvers in each section (lane changes, exits, platoon joins, etc.). The algorithm relies on a linear programming argument, which directly makes use of the spacing requirements discussed above.

3. AIR TRAFFIC MANAGEMENT SYSTEMS

Next generation ATM systems will include various automation tools, both on the ground and in the air (Pappas *et al.*, 1997). On the ground, the Center-TRACON Automation tool (CTAS), designed by NASA Ames, performs scheduling functions for air traffic flow coordinators and generates advisories to air traffic controllers in order to safely separate aircraft and guide them to their desired destinations. These advisories will be received by sophisticated on board Flight Management Systems (FMS) which may negotiate with controllers the aircraft's preferred alternatives. The goal of the FMS is, in a sense to provide the maximum number of safe alternatives, among which the scheduling algorithms can pick the optimal schedule.

[1] γ reflects the fact that the braking effort in general increases for followers further back in the platoon (Swaroop, 1994).

3.1 *CTAS Architecture*

The Air Traffic Control system consists of three types of ground control facilities: Centers which control en-route flights, Terminal Radar and Approach Control facilities (TRACON) which control arriving and departing flights close to airports, and Airport Control Towers which control traffic in the immediate vicinity of the airport. The Center-TRACON Automation System (Erzberger *et al.*, 1993) provides advisories for the air traffic controllers, in an attempt to increase airspace utilization, reduce delays, fuel consumption and controller workload as well as improve safety. CTAS consists of three main components:

- **Traffic Management Advisor (TMA)**
- **Descent Advisor (DA)**
- **Final Approach Spacing Tool (FAST)**

TMA and DA coexist and operate in Center airspace whereas FAST operates as a stand alone in TRACON airspace. The architecture of CTAS is shown in Figure 4. Thought of as a large input-output system, CTAS receives input from:
Controllers, who set the capacity and acceptance rates for various runways, airports and the TRACON.
Radar Daemon, which periodically receives aircraft state information regarding position, altitude, speed, aircraft type and flight plan.
Weather Daemon, which receives weather information from the National Weather service.

CTAS utilizes detailed databases which include aircraft, aerodynamic and engine models for all aircraft types. These models are used to perform accurate trajectory prediction for each aircraft. Other databases define the local airspace structure in terms of way-points and routes, and site adaptation data that define TRACON, airport and runway configurations and site specific constraints. After all calculations are performed, the main outputs of CTAS are:
Landing Schedules Given capacity, flow constraints and separation requirements, CTAS performs runway allocation and scheduling for all arriving aircraft.
Advisories In addition to scheduling, CTAS computes heading, altitude and speed advisories.

We now describe the operation of some CTAS tools that assist Center controllers, namely TMA and DA. For a more formal approach to the modeling and safety analysis of CTAS see (Lygeros *et al.*, 1998*b*).

3.2 *TMA and DA*

The goal of TMA is to assist the Traffic Management Coordinator (TMC) who monitors the

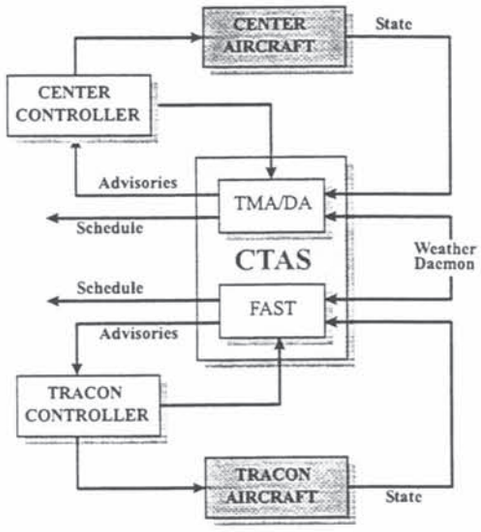


Fig. 4. CTAS Architecture

traffic flow. TMA therefore tries to minimize time delays while satisfying the various capacity and flow constraints but does not resolve conflicts between aircraft. TMA is therefore concerned with efficiency rather than safety issues.

The TMC may set certain capacity limits for the TRACON. At each computation cycle, TMA first obtains the state and flight plan of each aircraft from the Radar Daemon. The radar data is passed on to the Route Analyzer (RA) which calculates possible horizontal routes from the current state all the way to possible runways. Each route may be locally perturbed using predefined Degrees of Freedom (DOF). Degrees of freedom are site-specific, predefined path deviations which controllers use to resolve conflicts or to delay aircraft. For each aircraft, many possible routes, to various runways, each with many DOF perturbations are computed. For each route, RA uses the Trajectory Synthesis (TS) algorithm which computes a four dimensional profile. TS selects the appropriate aircraft aerodynamic and engine models from the database and uses them along with weather data from the weather daemon for trajectory prediction. The four dimensional profiles are used to calculate the Estimated Times of Arrival (ETAs) at various scheduling points including the TRACON metering gate and the runway threshold.

The possible TRACON and runway ETAs are used by the Dynamic Planner (DP), which tries to select from the possible ETAs the ones that minimize time delays and satisfy certain separation requirements. The DP produces a sequence of Scheduled Times of Arrival (STA) for all aircraft at both all TRACON gates and runways. Initial incarnations of DP were based on mathematical programming techniques resulting in the Implicit Enumeration algorithm (Brinton, 1992).

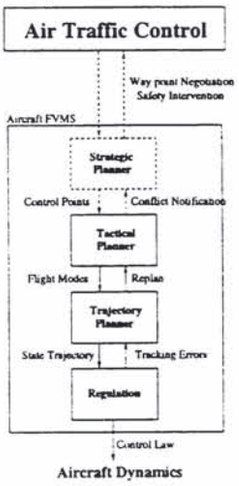


Fig. 5. FMS Architecture

Given the TMA produced time-line, the Descent Advisor (DA) tries to compute fuel efficient, conflict free descent advisories which try to meet the TMAs calculated schedule of entry to the TRACON. After TMA provides possible ETA calculations per route, DOF, and aircraft, and given a nominal sequence which meets the TRACON acceptance rates, the Profile Selector (PFS) selects routes and DOF for each aircraft such that the ETA matches the TMA produced STA and the routes are conflict free. In addition, it generates advisories to Center controllers in order to meet these objectives. These advisories are displayed on the controllers' graphical user interface which are then transmitted to the aircraft by voice.

3.3 *Flight Management Systems*

In next generation ATM systems, aircraft will be equipped with various planning and control algorithms. Each aircraft's FMS implements the hierarchical structure shown in Figure 5.

Strategic Planner: The main objectives of the strategic planner are to design a coarse, optimal trajectory for the aircraft and to resolve conflicts between aircraft. The trajectory may be adjusted in order to adapt to weather patterns, potential conflicts and airport traffic. The strategic planners of conflicting aircraft determine a sequence of maneuvers which will result in conflict-free trajectories (Tomlin *et al.*, 1998). The strategic planner sends its designed trajectory to the Tactical Planner in the form of a sequence of control points and/or a maneuver.

Tactical Planner: The tactical planner refines the strategic plan by interpolating the control points with smooth output trajectories. The tactical planner uses a simple kinematic model of the aircraft for all trajectory calculations. The output trajectory is then being passed to the Trajectory Planner.

Trajectory Planner: The trajectory planner detailed dynamic models of the aircraft, wind data, and the output trajectory of the Tactical Planner, to design full state and input trajectories for the aircraft, and a sequence of *flight modes* necessary to execute the dynamic plan. Flight modes represent different modes of operation of the aircraft and correspond to controlling different outputs in the aircraft dynamics. A derivation of the flight mode logic necessary for safe operation of a CTOL (Conventional Take Off and Landing) aircraft is presented in (Lygeros *et al.*, 1997). The task of the trajectory planner is complicated by the presence of non-minimum phase dynamics (Tomlin *et al.*, 1995) and actuator saturation (Pappas *et al.*, 1995).

Regulation Layer: Once a feasible dynamic trajectory has been determined, the regulation layer is asked to track it. In the presence of large external disturbances tracking can severely deteriorate in which case tracking errors are passed back to the trajectory planner, to facilitate re-planning if necessary.

4. CONCLUSIONS

In this paper, we have presented the large scale system paradigm through two applications of major interest: Automated Highway Systems (AHS) and Air Traffic Management (ATM) systems. In both cases, emphasis was placed on the interaction between higher level systems concerned with efficiency and lower level systems concerned with safety. Due to the complexity of the overall system, more elaborate hybrid techniques, such as those in (Lygeros *et al.*, 1997) are needed. In addition, higher levels of the hierarchy, utilize coarser system models or coarser abstractions. This raises the interesting notions of consistent abstractions or implementability, which is the ability of a lower level system to execute the commands of a higher level system. Initial results along this direction can be found in (Pappas *et al.*, 1998)

Acknowledgments: The authors would like to thank Datta Godbole and Claire Tomlin for their contributions in various sections of this manuscript.

5. REFERENCES

Brinton, C.R. (1992). An implicit enumeration algorithm for arrival aircraft scheduling. In: *Proceedings of the 11th Digital Avionics Systems Conference.* Seattle, WA.

Broucke, M. E. and P. Varaiya (1996). A theory of traffic flow in automated highway systems. *Transportation Research Part C* **4C**(4), 181–210.

Erzberger, H., T.J. Davis and S. Green (1993). Design of center-tracon automation system. In: *Proceedings of the AGARD Guidance and Control Syposium on Machine Intelligence in Air Traffic Management.* pp. 11.1–11.12.

Godbole, D. and J. Lygeros (1994). Longitudinal control of the lead car of a platoon. *IEEE Transactions on Vehicular Technology* **43**(4), 1125–1135.

Godbole, D. and J. Lygeros (1997). Tools for safety and throughput analysis of automated highway systems. In: *American Control Conference.* pp. 2031–2035.

Lygeros, J., C. Tomlin and S. Sastry (1997). Multi-objective hybrid controller synthesis. In: *Proceedings of HART97* (Oded Maler, Ed.). pp. 109–123. Number 1201 In: *LNCS.* Springer Verlag. Berlin.

Lygeros, J., D. Godbole and S. Sastry (1998*a*). Verified hybrid controllers for automated vehicles. *IEEE Transactions on Automatic Control* **43**(4), 522–539.

Lygeros, J., G. Pappas and S. Sastry (1998*b*). An approach to the verification of the Center-TRACON automation system. In: *Hybrid Systems: Computation and Control*, number 1386 In: *LNCS.* Springer Verlag. pp. 289–305.

Pappas, G., C. Tomlin, J. Lygeros, D. Godbole and Shankar Sastry (1997). A next generation architecture for air traffic management systems. In: *Proceedings of the 1997 IEEE Conference in Decision and Control.* pp. 2405–2410.

Pappas, G., G. Lafferriere and S. Sastry (1998). Hierarchically consistent control systems. Technical Report UCB/ERL M98/16. University of California, Berkeley, CA.

Pappas, G., J. Lygeros and D. Godbole (1995). Stabilization and tracking of feedback linearizable systems under input constraints. In: *IEEE Control and Decision Conference.* pp. 596–601.

Swaroop, D. (1994). String Stability of Interconnected systems: an application to platooning in automated highway systems. PhD thesis. Department of Mechanical Engineering, University of California, Berkeley.

Tomlin, C., G. Pappas and S. Sastry (1998). Conflict resolution for air traffic management: a case study in multi-agent hybrid systems. *IEEE Transactions on Automatic Control* **43**(4), 509–521.

Tomlin, C., J. Lygeros, L. Benvenuti and S. Sastry (1995). Output tracking for a non-minimum phase dynamic CTOL aircraft model. In: *IEEE Control and Decision Conference.* pp. 1867–1872.

Varaiya, P. (1993). Smart cars on smart roads: problems of control. *IEEE Transactions on Automatic Control* **AC-38**(2), 195–207.