COOPERATIVE ATTITUDE SYNCHRONIZATION IN SATELLITE SWARMS: A CONSENSUS APPROACH

Alain Sarlette (alain.sarlette@ulg.ac.be) [2]
Rodolphe Sepulchre (r.sepulchre@ulg.ac.be)

Systems and Modeling, Montefiore Institute (B28), University of Liège, 4000 Liège, Belgium

Naomi Ehrich Leonard (naomi@princeton.edu)

MAE Department, Engineering Quadrangle, Princeton University, NJ 08544, USA

Abstract: The present paper considers the problem of autonomous synchronization of attitudes in a swarm of spacecraft. Building upon our recent results on consensus on manifolds, we model the spacecraft as particles on $SO(3)$ and drive these particles to a common point in $SO(3)$. Unlike the Euler angle or quaternion descriptions, this model suffers no singularities nor double-points. Our approach is fully cooperative and autonomous: we use no leader nor external reference. We present two types of control laws, in terms of applied control torques, that globally drive the swarm towards attitude synchronization: one that requires tree-like or all-to-all inter-satellite communication (most efficient) and one that works with nearly arbitrary communication (most robust).

Keywords: Consensus algorithms, Attitude control, Swarm control, Formation flight, Satellite swarms, Autonomous swarms, Synchronization

1. INTRODUCTION

Many modern space mission concepts involve the use of multiple satellites flying in formation. Mostly, the objective is to get (virtual) structures in orbit that are substantially larger than what current launch technologies can directly handle. Potential applications arising in current studies include resolution enhancement through multiple-spacecraft SAR (the InSAR concept, or ONERA's Romulus study), interferometry (ESA's Darwin project, NASA's equivalent Terrestrial Planet Finder project or NASA's Constellation-X project) or supersized focal length (ESA's XEUS

[1] This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office. The scientific responsibility rests with its authors.
[2] Corresponding author. Supported as an FNRS fellow (Belgian Fund for Scientific Research). This work is achieved while the first author is a Visiting Student Research Collaborator at the MAE department of Princeton University, whose hospitality and intellectually inspiring environment are gratefully acknowledged. Current address: MAE Department, Engineering Quadrangle, Princeton University, NJ 08544, USA. Financial support for this visit is mainly provided through the 1st Odissea prize 2005 initiated by the Belgian Senate.

project, derived from the Symbol-X concept of CNES), sensitivity increasing through screens on secondary spacecraft (the American New World Discoverer concept for the JWST) or large-scale measurements (the ESA-NASA cooperative mission LISA), and autonomous in-orbit assembly of large real structures (projects are still at a draft level; see Izzo and Pettazzi [2005] and McInnes [1996] for example). Other advantages of spacecraft formations are their robustness with respect to single spacecraft failure, and the reconfigurability of the swarm to fit specific observation requirements.

A central problem in formation flight control is to ensure proper synchronization of the spacecraft, *i.e.* to bring them to and keep them in the desired formation. Many interesting studies are devoted to *position* synchronization; a thorough survey would require a longer discussion and we refer the interested reader to citations in the references of the present paper.

Attitude synchronization has attracted somewhat less attention. Algorithms that asymptotically stabilize synchronized satellite attitudes are presented in Lawton and Beard [2002] and VanDyke and Hall [2004]. Interconnections among satellites are limited, and convergence is proved for a behavioral algorithm combining tracking of a desired attitude, eigenaxis rotations and synchronization of the swarm. However, these results strongly depend on the tracking of a common external reference: when the latter is suppressed, the limited attraction region for which synchronization is guaranteed vanishes to the empty set. In Bondhus et al. [2005] and Krogstad and Gravdahl [2006], attitude synchronization is considered with a leader/follower approach. In that case, the leader spacecraft can be seen as a reference which is tracked by the followers. Robustness of this approach strongly relies on the reliability of the leader spacecraft and on the ability of all the followers to track it. Control algorithms are presented that globally stabilize attitude synchronization, but the use of the convenient but non-unique quaternion representation of rigid-body orientations produces unwanted artifacts in the satellites' motions: sometimes a satellite that has an attitude very close to the leader moves in the opposite direction to come back from another side. It seems that quaternions are absolutely reliable as long as relative orientations are considered individually, but can run into problems when several orientations are combined without a common external reference. Nair and Leonard [2004] consider the attitude synchronization problem without external reference on $SO(3)$. In fact, their artificial coupling potentials use the same distance measure as we do in the present work. Using the Method of Controlled Lagrangians, local stability of a synchronized state is studied and achieved in a specific situation (final synchronized rotation around the short principal axis, the satellites interconnected in a path topology). In addition to being local, this result is *not asymptotic*, meaning that the satellites remain close to the equilibrium but do not converge towards it.

Our goal is to provide control laws that drive the swarm towards attitude synchronization from any initial configuration, without using any leader or external reference, and for various satellite interconnection topologies. We also formulate the problem directly on $SO(3)$ to avoid problems related to the Euler angle and quaternion representations.

Building on our recent results on consensus on manifolds, the proposed approach considers attitude synchronization as a consensus problem on $SO(3)$. As in the traditional consensus approach, we start by generating desired trajectories on a first-order, kinematic model. In a second step, the related satellite dynamics are incorporated to derive control laws in terms of torques.

Since our approach focuses on states that are far away from the desired equilibrium and inherently includes strong robustness considerations, the resulting control laws might most probably be useful for initial deployment of the formation or for recovery after strong transient perturbations. In general, most theoretical studies are far from final science operations requirements, where the key issue for navigation and control is accuracy; see Beugnon et al. [2005] for an example of a practical GNC implementation for the Darwin mission.

The paper is organized as follows. In Section 2, we review the consensus strategy in the specific setting of synchronization on $SO(3)$. In Section 3, two types of kinematic consensus algorithms are presented. The first one just requires relative attitude sensing, but only retains its *global* stabilization property for specific requires fixed and bidirectional satellite interconnections. The second one requires explicit inter-satellite communication but assuredly globally converges for a very large class of interconnection topologies, including varying ones. Section 4 considers the dynamical implementation in order to asymptotically track the kinematic trajectories defined in Section 3. Classical tracking approaches (computed torque and high gain control) lead to some specific "consensus tracking" control torques.

2. CONSENSUS ON $SO(3)$

2.1 The special orthogonal group

The special orthogonal group $SO(3)$ is the natural state space to describe rigid-body orientations in $\mathbb{R}^3$; it can equivalently be defined as the group of rotations or as the group of positively oriented orthonormal bases in $\mathbb{R}^3$. $SO(3)$ is a compact, connected and homogeneous manifold of dimension 3.

In its canonical representation, a point q of $SO(3)$ is characterized by a real 3×3 orthogonal matrix Q with determinant equal to $+1$; in this paper, we represent the position of an agent k on $SO(3)$ by such a matrix Q_k. In this representation, the inverse of a group element is simply the inverse or transpose matrix and the group action is represented by matrix multiplication. Hence if Q_1 and $Q_2 \in SO(3)$ are the positions of two elements with respect to an inertial frame, then $Q_1^T Q_2 \in SO(3)$ represents the relative position of 2 with respect to 1. The gradient $D_q[F](q_1)$ of a function F along $SO(3)$ at a point q_1 is simply the projection of the gradient $D_Q[F](q_1)$ in $\mathbb{R}^{3\times3}$ onto the tangent space to $SO(3)$ at q_1. The representation of the tangent space at the identity I_3 is the set of skew-symmetric matrices. Using group operations, we deduce

$$D_q[F](q_1) = Q_1 \frac{Q_1^T D_Q[F](q_1) - D_Q[F](q_1)^T Q_1}{2} \; .$$

The following lemmata will be useful; their proofs can be found in Sarlette and Sepulchre [2006]. Just a note about the *polar decomposition*. Any matrix $B \in \mathbb{R}^{3\times3}$ can be decomposed into a product UR where U is orthogonal and R is symmetric positive semi-definite. R is always unique, U is unique if B is non-singular. If B has rank 2, U is unique in $SO(3)$.

Lemma 1: *The matrix $Q^T B$ with $Q \in SO(3)$ and $B \in \mathbb{R}^{3\times3}$ is symmetric iff $Q = UHJH^T$ where*

$$J = \begin{pmatrix} -I_l & 0 \\ 0 & I_{3-l} \end{pmatrix} ,$$

$B = UR$ is a polar decomposition of B, l is even if $\det(U) > 0$ and odd if $\det(U) < 0$, and the columns of H contain (orthonormalized) eigenvectors of R.

Lemma 2: *For any $B \in \mathbb{R}^{3\times3}$, the local maxima of a linear function $F(Q) = \operatorname{trace}(Q^T B)$ are always global maxima as well.*

2.2 Cost function

The objective of consensus is basically to reach agreement in a set of agents - here on the attitude to adopt (pardon the pun). Mathematically, this requires defining a measure of disagreement between the agents which is then driven to zero. In the present case, an obvious measure of disagreement is the set of all pairwise distances between agents on $SO(3)$. Computing the inherent Riemannian distance $d(q_1, q_2)$ between two points q_1 and q_2 along $SO(3)$ is a complex task. Instead, we choose to approximate this distance by the Euclidean distance $d(Q_1, Q_2)$ between the matrices Q_1 and $Q_2 \in SO(3)$ in the embedding space $\mathbb{R}^{3\times3}$:

$$\begin{aligned} d(Q_1, Q_2) &= \|Q_1 - Q_2\|_F \\ &= \sqrt{\operatorname{trace}\left((Q_1 - Q_2)^T (Q_1 - Q_2)\right)} \; . \end{aligned}$$

This is strictly similar to approximating the distance between two points along the circle $S^1 \cong SO(2)$ by the length of the chord between these points. For that reason, this approximation is sometimes called the *chordal distance*, as introduced by Conway et al. [1996] for the Grassmann manifold. Thanks to the finite curvature of $SO(3)$, the chordal distance and the Riemannian distance are asymptotically equivalent when they converge to zero; in particular, $d(q_1, q_2) = 0 \Leftrightarrow d(Q_1, Q_2) = 0$ and $d(q_1, q_3) > d(q_1, q_2) \Leftrightarrow d(Q_1, Q_3) > d(Q_1, Q_2)$. Our main motivation for choosing the chordal distance is that, unlike for the Riemannian distance, the derivative $D_{q_1}[d^2(Q_1, Q_2)](q_1, q_2)$ has no discontinuities. In practice, the chordal distance has the advantage of being easily computable: its square is simply the sum of the squared elements of the matrix $Q_1 - Q_2$. In Moakher [2002], this distance is used to define the *projected arithmetic mean* of N particles on $SO(3)$ as the (set of) point(s) Q that minimize(s) the sum of the squared chordal distances from the point Q to all the particles. It is shown that this mean is simply the orthogonal projection on $SO(3)$ of the *Euclidean centroid* C_e of the particle positions in $\mathbb{R}^{3\times3}$,

$$C_e = \frac{1}{N} \sum_{k=1}^{N} Q_k \; .$$

Moakher [2002] also shows that the projected arithmetic mean and the *Karcher mean*, which uses the inherent Riemannian distance, are strictly equivalent for $N = 2$.

To obtain a scalar disagreement measure, it is natural to take some weighted sum of pairwise disagreements. This sum needs not to include all terms to possess the synchronized state as its global minimum. Explicitly, consider that the satellites are interconnected according to a directed graph G and denote by $j \rightsquigarrow k$ the fact that j sends its relative attitude to k (or k measures the relative attitude of j). For later reference, we also define the undirected graph associated to a directed graph G to be equivalent to G with directions discarded. We define the *consensus cost function*

$$P_G(Q_1...Q_N) = \frac{1}{2N} \sum_{k=1}^{N} \sum_{j \rightsquigarrow k} (d(Q_j, Q_k))^2 \ . \quad (1)$$

Since $Q_k^T Q_k = I_3$, this can be rewritten as

$$P_G(Q_1...Q_N) = \frac{1}{N} \sum_{k=1}^{N} \sum_{j \rightsquigarrow k} 3 - \text{trace}(Q_j^T Q_k) \ .$$

For the particular case where each satellite is connected to all the others, we drop the index G; in that case we may rewrite (1) as

$$P(Q_1...Q_N) = N(3 - \|C_e\|^2) \ . \quad (2)$$

The second term in (2) is analogous to the order parameter used to study synchronization on the circle, most famous in the framework of the Kuramoto model. See Scardovi et al. [2007] for consensus algorithms on the circle similar to those of Section 3 in the present paper.

2.3 Synchronization strategy

It is obvious that the synchronized state is the unique configuration that globally minimizes (1) whenever the satellites are weakly interconnected (*i.e.* when you can find a path from any satellite to any other satellite by travelling along interconnections only, but discarding their orientation). This suggests to simply apply consensus cost function minimization in order to drive the swarm towards synchronization. However, because of the manifold structure of $SO(3)$, this is not always a convex problem; unlike in Euclidean spaces, the consensus cost function could have local minima. This is the fundamental reason why algorithmically, global attitude synchronization is a much more difficult task than position synchronization. The following proposition identifies two situations where there are no local minima.

Proposition 1: *If G, the* undirected graph associated to *the satellite interconnections is a complete graph or a tree, then attitude synchronization is the only minimum of P_G.*

Proof: At a minimum of P_G, each satellite must be at a local minimum for its own movements.

For the complete graph, consider P with all agents fixed except k. Q_k must be a maximizer of $\text{trace}(Q_k^T \sum_{j \neq k} Q_j)$; according to Lemma 2, there are only global maxima. Furthermore, it is obvious that $\text{trace}(Q_k^T Q_k) > \text{trace}(Y^T Q_k)$ whenever $Y \neq Q_k \in SO(3)$. As a consequence, Q_k is at the only maximizer of $\text{trace}(Y^T C_e)$ over $Y \in SO(3)$; since this must hold $\forall k$, all the satellites must be synchronized at the same attitude.

For the tree, start with all agents fixed except a leaf k. The obvious unique global maximum of $P_G(Q_k)$ occurs when k is synchronized with its parent. Now consider synchronized moves of a parent j and its leaves. Variations of P_G only involve the link between j and its own parent, which brings us back to the previous situation. An inductive argument is then used up to the root.

△

Note that the much "cheaper" tree interconnection has a similar disadvantage as the leader-follower approach: if a satellite fails, all its children and further dependence are disconnected from the rest of the swarm.

3. CONSENSUS ALGORITHMS

3.1 Gradient control for undirected graphs

Gradient control leads to descent algorithms for P_G. Explicitly, we write

$$Q_k^T \dot{Q}_k = -\alpha_k \, Q_k^T D_q[P_G](q_k)$$
$$= \alpha_k \, (C_k - C_k^T) \quad (3)$$
$$\text{with } C_k = \frac{1}{2N} \left(\sum_{j \rightsquigarrow k} Q_k^T Q_j + \sum_{k \rightsquigarrow j} Q_k^T Q_j \right) . \quad (4)$$

The left member of (3) is the attitude variation of k with respect to an inertial frame, expressed in the reference frame of k. Note that the coefficients $\alpha_k(t) \geq \overline{\alpha} > 0$ may depend on time and be chosen independently for each satellite. This leaves some freedom for adaptation to practical constraints.

According to the definition (4) of C_k, each satellite k must know the relative orientation of the satellites to which it is connected *in either direction.* This only makes sense if the interconnection graph is undirected, *i.e.* k sends information to j iff it gets information from j. The following proposition is obvious from Proposition 1 and the properties of gradient systems.

Proposition 2: *If the graph corresponding to the satellite interconnections is an undirected tree or complete graph and each satellite orients itself according to (3), then the only stable configuration for the swarm is attitude synchronization.*

Note that according to the discussion of the previous section, attitude synchronization is still a locally stable equilibrium of (3) for other undirected graphs, though other local equilibria may exist for some of them. In fact, simulations with directed (dropping the second term of (4)) and randomly time-varying graphs, modeling extreme conditions of link failures, invariably converge to synchronized attitudes, suggesting the strong robustness of algorithm (3).

3.2 Auxiliary variables for general graphs

The goal of this second approach is to derive an algorithm whose convergence is theoretically proven when the satellite interconnections are varying and/or directed. A strong convergence result has been obtained for Euclidean spaces by Moreau [2004]. The precise assumption on the interconnections is *uniform connectedness.* A time-varying graph $G(t)$ is called uniformly connected if there exist a node k and time horizons δ and T such that $\forall t$, there is a directed path from k to all other nodes, in the graph that contains all the interconnections appearing for at least δ seconds in $[t, t+T]$. This is a very weak condition on the satellite interconnections. The argument of Moreau [2004], based on convexity, cannot be extended to $SO(3)$. However, we can use it to achieve consensus on auxiliary variables X_k in $\mathbb{R}^{3\times3}$. We then drive the satellite attitudes to the projection of the consensus value X_∞ on $SO(3)$. One easily verifies that the latter corresponds to the maximizer of $\text{trace}(Q^T X_\infty)$ over $Q \in SO(3)$. As a consequence of Lemmata 1 and 2, it can be computed by local maximization and it is unique whenever the smallest eigenvalue of X_∞ has multiplicity 1. By translational invariance, the latter condition is ensured with probability 1 for randomly chosen $X_k(0)$. This leads to the algorithm

$$\tfrac{d}{dt}(Q_k^T X_k) = \beta_k \sum_{j \rightsquigarrow k} (Q_k^T Q_j)(Q_j^T X_j) + (Q_k^T \dot{Q}_k)^T (Q_k^T X_k) \quad (5)$$

$$Q_k^T \dot{Q}_k = \alpha_k\, (Q_k^T X_k - X_k^T Q_k) \quad (6)$$

with $\beta_k(t), \alpha_k(t) > 0$. Again, we chose to express the variables in the body frames of the satellites. This is important because (5) absolutely requires inter-satellite communications. When the relative attitudes $Q_k^T Q_j$ are measured as geometric quantities, the values $Q_j^T X_j$ are communicated as sets of scalars.

Proposition 3 : *If the graph $G(t)$ associated to the satellite interconnections is piecewise continuous and uniformly connected, then the only stable equilibrium of (5),(6) is attitude synchronization.*

Proof: In an inertial frame, (5) is simply

$$\dot{X}_k = \beta_k \sum_{j \rightsquigarrow k} X_j \quad (7)$$

Without going into details, the result of Moreau [2004] ensures that the independent system (7) exponentially converges to a consensus value X_∞ under the given conditions. Therefore X_k may asymptotically be replaced by X_∞ in (6), which then converges to the projection of X_∞ on $SO(3)$; the latter is a single point except for the $X_k(0)$ in a set of measure zero in $(\mathbb{R}^{3\times3})^N$.

$\triangle$

4. DYNAMIC IMPLEMENTATION: CONSENSUS TRACKING

The dynamics of satellite orientations are described by the equations

$$J\dot{\omega}_k = J\omega_k \times \omega_k + \tau_k \quad (8)$$

$$Q_k^T \dot{Q}_k = [\omega_k]^\wedge \quad (9)$$

where $J = \text{diag}(J_1, J_2, J_3)$, with $J_1 > J_2 > J_3$, is the moment of inertia matrix in (principal) body frame, ω_k and τ_k are the angular velocity and control torque respectively expressed in body frame, and the operator $[\cdot]^\wedge : \mathbb{R}^3 \to \mathbb{R}^{3\times3}$ denotes the skew-symmetric matrix implementing the cross-product: $[x]^\wedge y = x \times y\ \forall y$. For $x = (x_1\ x_2\ x_3)^T$ this implies

$$[x]^\wedge = \begin{pmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{pmatrix}.$$

The inverse of $[\cdot]^\wedge$ is denoted $[\cdot]^\vee : \mathbb{R}^{3\times3} \to \mathbb{R}^3$. Unlike Newton's law for position control, equations (8)-(9) for orientation control are nontrivial, even nonlinear, in the absence of external torques.

The consensus algorithms of Section 3 define desired angular velocities $\omega_k^{(d)} = [Q_k^T \dot{Q}_k]^\vee$ for each satellite based on the current positions of themselves and their neighbors. It remains to design τ_k to ensure that the satellites individually track these desired velocity fields. In fact, a simpler control strategy can be based on the approach of Section 3.2, since a desired orientation $Q_k^{(d)}$ can be directly defined as the projection of X_k on $SO(3)$. In this case, it may even be useless to move the satellites before the X_k have converged to a consensus value X_∞ by applying (5) only. Subsequently moving the agents individually to $Q_\infty^{(d)}$ just requires a global attitude *stabilization* controller instead of tracking capabilities. Attitude tracking algorithms abound in the literature. The following applies two classical approaches of control theory to tracking of $\omega_k^{(d)}$: "computed torque" and "high gain". Proofs for the coupled system are omitted because of space limitations.

We want to impose an exponential convergence of ω_k towards $\omega_k^{(d)}$:

$$\tfrac{d}{dt}(\omega_k - \omega_k^{(d)}) = -\gamma_k(\omega_k - \omega_k^{(d)}) \quad (10)$$

where $\gamma_k > 0$. The "computed torque" method simply implements this law on the dynamical system by including a control torque that exactly cancels the free rigid body dynamics. This yields

$$\tau_k = J \tfrac{d}{dt}\omega_k^{(d)} - \gamma_k(\omega_k - \omega_k^{(d)}) - (J\omega_k) \times \omega_k \,. \quad (11)$$

Proposition 4: *Consider the control torque (11), where $\omega_k^{(d)}$ is defined by a consensus algorithm of Section 3. If the interconnection graph is uniformly connected and (5)-(6) is used, then the only stable limit set is synchronization of the orientations with $\omega_k = 0\ \forall k$. If the interconnection graph is fixed and undirected and (3) is used, the stable limit set consists of the local minima of P_G with $\omega_k = 0\ \forall k$.*

Some adaptation law should be added in order to ensure that the natural dynamics are perfectly cancelled. Alternatively, "high gain" tries to dominate the natural dynamics by imposing a strong enough control gain γ_k. This reduces the control torque to

$$\tau_k = J \tfrac{d}{dt}\omega_k^{(d)} - \gamma_k(\omega_k - \omega_k^{(d)}) \,. \quad (12)$$

Proposition 5: *Replacing (11) by (12), the results of Proposition 4 are recovered at least when*

$\gamma_k > \alpha_k \frac{3J_1}{\sqrt{2}J_3} \|X_{max}(0)\|_F$ *when using (5)-(6),*

$\gamma_k > \alpha_k \frac{\sqrt{6}J_1}{J_3} d_k$ *when using (3),*

where $\|X_{max}(0)\|_F$ is the maximal Frobenius norm of the $X_k(0)$ and d_k is the number of satellites sending information to k.

5. CONCLUSION

The consensus approach to satellite attitude synchronization leads to particularly robust control algorithms. It also leaves sufficient flexibility for practical implementations. Indeed, many tracking controllers can be used in place of the theoretically simple ones considered in Section 4. The real control performance depends on the choice of the tracking controller.

The satellites have to measure their own absolute angular velocities (gyroscopic sensor) in order to implement the "consensus tracking" control laws of the present paper. In Sarlette et al. [2007], we discuss an alternative synchronization strategy, based on Nair and Leonard [2004], which works (though only locally) without absolute angular velocity measurements. Taking an inherently dynamical approach, this strategy also avoids to explicitly counter the free rigid body dynamics as is done in the present paper.

REFERENCES

For information on the space mission concepts, see the corresponding organizations' web pages.

C. Beugnon, E. Buvat, M. Kersten, and S. Boulade. GNC design for the DARWIN spaceborne interferometer. *Proc. 6th ESA Conf. on Guidance, Navigation and Control Systems*, 2005.

A.K. Bondhus, K.Y. Pettersen, and J.T. Gravdahl. Leader/follower synchronization of satellite attitude without angular velocity measurments. *Proc. 44th IEEE Conf. on Decision and Control*, pages 7270–7277, 2005.

J.H. Conway, R.H. Hardin, and N.J.A. Sloane. Packing lines, planes, etc.: Packings in Grassmannian spaces. *Exper.Math.*, 5(2):139–159, 1996.

D. Izzo and L. Pettazzi. Equilibrium shaping: distributed motion planning for satellite swarm. *Proc. 8th Intern. Symp. on Artificial Intelligence, Robotics and Automation in space*, 2005.

T.R. Krogstad and J.T. Gravdahl. Coordinated attitude control of satellites in formation. In *Group Coordination and Cooperative Control*, volume 336 of *Lecture Notes in Control and Information Sciences*, chapter 9, pages 153–170. Springer-Verlag, 2006.

J.R. Lawton and R.W. Beard. Synchronized multiple spacecraft rotations. *Automatica*, 38: 1359–1364, 2002.

C.R. McInnes. Distributed control for on-orbit assembly. *Advances in the Astronautical Sciences*, 90:2079–2092, 1996.

M. Moakher. Means and averaging in the group of rotations. *SIAM J. Matrix Anal. Appl.*, 24 (1):1–16, 2002.

L. Moreau. Stability of continuous-time distributed consensus algorithms. *Proc. 43rd IEEE Conf. on Decision and Control*, pages 3998–4003, 2004.

S. Nair and N.E. Leonard. Stabilization of a coordinated network of rotating rigid bodies. *Proc. 43rd IEEE Conf. on Decision and Control*, pages 4690–4695, 2004.

A. Sarlette and R. Sepulchre. Consensus optimization on manifolds. *Submitted to SICON special issue on Control and Optimization in Cooperative Networks*, 2006.

A. Sarlette, R. Sepulchre, and N.E. Leonard. Autonomous rigid body attitude synchronization. *Submitted to the 46th IEEE Conf. on Decision and Control*, 2007.

L. Scardovi, A. Sarlette, and R. Sepulchre. Synchronization and balancing on the N-torus. *Systems and Control Letters*, 56(5):335–341, 2007.

M.C. VanDyke and C.D. Hall (advisor). Decentralized coordinated attitude control of a formation of spacecraft. *Master Thesis, VirginiaTech*, 2004.