

PARTICLE FLOW SMC DELTA-GLMB FILTER

Augustin-Alexandru Saucan, Yunpeng Li and Mark Coates

McGill University, Department of Electrical and Computer Engineering
Montreal, QC, Canada (augustin.saucan@mail.mcgill.ca)

ABSTRACT

In this paper we derive a particle flow particle filter implementation of the δ -GLMB filter. The particle flow is employed to produce a measurement-driven importance distribution that serves as a proposal in the δ -GLMB particle filter. Flow-induced proposals are developed for both types of targets encountered in the δ -GLMB filter, i.e., persistent and birth targets. Numerical simulations reflect the improved performance of the proposed filter as opposed to classical bootstrap implementations.

Index Terms— target tracking, random finite set, Bayesian estimation, particle filter, particle flow.

1. INTRODUCTION

Multi-target tracking addresses the problem of estimating the trajectories of an unknown and time-varying number of targets from a sequence of observations. Labeled Random Finite Sets (RFS) have been recently proposed in [1] for the specific task of trajectory inference. Most notably, the δ -GLMB family of labeled RFS densities was shown to be closed under the Chapman-Kolmogorov and Bayes equations. The resulting δ -GLMB filter was proposed in [2] alongside with Gaussian mixture and Sequential Monte Carlo (SMC) implementations.

For state systems with highly informative measurements or with a high-dimensional state space, the Bootstrap Particle Filter (BPF) suffers from the phenomenon of weight degeneracy [3]. Additionally, we show in this paper that the BPF δ -GLMB exhibits the same problems whenever the measurement likelihood is concentrated as opposed to the prior. Particle flow [4] was proposed as a solution to single target filtering in such situations. Indeed, particle flow links the prior and posterior distributions through a log-homotopy and uses partial differential equations to migrate prior particles toward the posterior distribution. In [5] and [6], particle flow filters are developed to generate a proposal distribution which serves as an importance distribution in a particle filter context. The particle flow proposal takes into account the current measurement and is capable of generating samples in regions where

the posterior is significant. Most notable, in [6] the flow proposal and the prior distributions are linked through a bijective transformation and hence the importance weights are easily calculated.

In this paper, we propose a particle flow SMC implementation of the δ -GLMB filter. We develop the importance distributions and derive the importance weights for the two types of targets (birth and persistent) encountered in the δ -GLMB filter. The particle flow δ -GLMB retains an efficient parallelizable implementation similar to the standard δ -GLMB filter. The simulations conducted showcase the improved performance of the particle flow SMC δ -GLMB filter as opposed to the BPF δ -GLMB with reduced computational overhead.

The paper is structured as follows. Section 2 presents a brief background on labeled RFS filtering and the δ -GLMB density. Section 3 presents the proposed filter. Numerical results are presented in Section 4 and concluding remarks are given in Section 5.

2. BACKGROUND

In order to incorporate target tracks into the Bayes multi-target filtering framework, targets are identified by a label. The state of a labeled target, denoted with $\mathbf{x} = (x, l)$, is comprised of the target state $x \in \mathbb{X}$ and label $l \in \mathbb{L}$. The multi-target state $\mathbf{X}_k = \{\mathbf{x}_{k,1}, \dots, \mathbf{x}_{k,N(k)}\}$ is modeled as a Random Finite Set (RFS) [7], where the number of targets as well as the individual target states are random.

The objective is the estimation of the labeled multi-target posterior density $\pi_k(\mathbf{X}_k|Z_{0:k})$, which captures all information of the target states and labels given the measurement history $Z_{0:k} = \{Z_0, \dots, Z_k\}$. The multi-target Bayes filter [7] propagates π_k in time according to

$$\pi_k(\mathbf{X}_k|Z_{0:k}) = \frac{\mathbf{g}(Z_k|\mathbf{X}_k)\pi_{k|k-1}(\mathbf{X}_k)}{\int \mathbf{g}(Z_k|\mathbf{X})\pi_{k|k-1}(\mathbf{X})\delta\mathbf{X}} \quad (1)$$

$$\pi_{k+1|k}(\mathbf{X}_k) = \int \mathbf{f}(\mathbf{X}_{k+1}|\mathbf{X}_k)\pi_k(\mathbf{X}_k|Z_{0:k})\delta\mathbf{X}_k, \quad (2)$$

where $\mathbf{g}(Z_k|\mathbf{X}_k)$ is the RFS counterpart of the single target likelihood function $g(z_k|x_k, l)$, and $\mathbf{f}(\mathbf{X}_{k+1}|\mathbf{X}_k)$ is the RFS counterpart of the single target transition kernel $f_{k+1|k}(x_{k+1}|x_k, l)$. The integral is a set integral defined in [8].

This work was supported by PWGSC contract W7707-145675/001/HAL funded by Defence R&D Canada.

The δ -GLMB family of multi-target densities, proposed in [1, 2], represents a closed form solution to the Chapman-Kolmogorov and the Bayes multi-target filtering equations of (1) and (2). A δ -GLMB density can be expressed as

$$\pi_k(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{(I, \xi) \in \mathcal{F}(\mathbb{L}_{0:k}) \times \Xi} w_k^{(I, \xi)} \delta_I(\mathcal{L}(\mathbf{X})) \left[p_k^{(\xi)} \right]^{\mathbf{X}}, \quad (3)$$

where $\mathcal{L} : \mathbb{X} \times \mathbb{L} \rightarrow \mathbb{L}$ is the projection $\mathcal{L}((x, l)) = l$ and $\Delta(\mathbf{X}) \triangleq \delta_{|\mathbf{X}|}(\mathcal{L}(\mathbf{X}))$ is the distinct label indicator. We employ the following notation $h^{\mathbf{X}} \triangleq \prod_{\mathbf{x} \in \mathbf{X}} h(\mathbf{x})$ where $\delta_{\mathbf{X}}(\mathbf{Y})$ is a generalized Kronecker delta which is equal to unity only when $\mathbf{X} = \mathbf{Y}$.

The sum in (3) is conducted over the pairs (I, ξ) , called hypotheses, which are composed of a label set I and a mapping $\xi : I \rightarrow Z_{0:k}$ that associates the measurements from Z or the void measurement (in case of miss detections) with the labels in I . Furthermore, note that each hypothesis in (3) has an associated weight $w_k^{(I, \xi)}$ and a set of target probability densities $p_k^{(x, l)}$ with $l \in I$.

The propagation of the δ -GLMB is achieved through prediction and update and involves the propagation of each hypotheses (I, ξ) . Both Gaussian mixture and SMC implementations were given in [2] and in the following section we give a particle flow implementation of the SMC δ -GLMB.

3. PARTICLE FLOW SMC δ -GLMB

In this section we give a brief description of the particle flow particle filter followed by its application to the δ -GLMB filter. Since within the δ -GLMB framework both persistent, i.e. targets surviving from the last time step and newly born targets coexist, their particle flows are addressed separately in Section 3.3 and Section 3.2. Practical implementation issues are discussed in Section 3.4.

3.1. Particle flow particle filter

At time step k and for each target with label l , the particle flow particle filter (PF-PF) [6] first propagates the i -th particle using the kinematic model: $\eta_0^i \sim f_{k|k-1}(x_k | x_{k-1}^i)$. An invertible mapping $\eta_1^i = T^i(\eta_0^i)$ then transports η_0^i via particle flow into regions where the posterior density is significant. The transformation is a continuous deformation indexed by the pseudo-time $\lambda \in [0, 1]$. The trajectory of η_λ^i at $\lambda \in [0, 1]$ is defined by an ordinary differential equation (ODE)

$$\frac{d\eta_\lambda^i}{d\lambda} = \zeta(\eta_\lambda^i, \lambda). \quad (4)$$

In practice the deformation is achieved by a finite series of transformations, that is, by the discretized integration of the above ODE with initial value η_0^i . Two numerical solutions exit for $\zeta(\eta_\lambda^i, \lambda)$ and are referred to as the exact Daum and Huang (EDH) [9] and the localized Daum and Huang

(LEDH) [10] filters. Both filters are applicable to Gaussian systems with non-linear measurement equations.

The authors of [6] prove that under mild assumptions on the smoothness of the measurement function, an invertible mapping is achievable between η_1^i and η_0^i and the importance distribution becomes

$$q(\eta_1^i | x_{k-1}^i, z_k) = \frac{p(\eta_0^i | x_{k-1}^i)}{|\det(\dot{T}^i(\eta_0^i))|}, \quad (5)$$

where $|\det(\dot{T}^i(\eta_0^i))|$ represents the absolute value of the Jacobian determinant of the mapping $T^i(\cdot)$. In the EDH version of the PF-PF, the evaluation of the determinant is avoided since the same mapping is applied to all particles while in the LEDH version of the PF-PF, the evaluation is straightforward since the mapping is a series of affine transformations. We refer the readers to [6] for further implementation details for both EDH and LEDH variants of the PF-PF.

The δ -GLMB filter has different propagation schemes for the persistent and birthed targets. Therefore, the design of the invertible flow will be treated separately in the following for the two cases. For both types of targets, we will only consider the case where a detected measurement is available since in the case of a missed detection no measurement-dependent importance distribution can be constructed.

3.2. δ -GLMB birth flow

The δ -GLMB updates with the measurement z a birth target having the prior $p_b(x, l)$ via the equation

$$p_b(x, l | z) \propto p_D(x, l) g(z | x, l) p_b(x, l). \quad (6)$$

Let $\{\eta_0^i\}_{i=1}^N$ be a set of particles sampled from the prior $p_b(x, l)$ and $\{\eta_1^i\}_{i=1}^N$ be the invertible-flow transformation of these particles. Note that the label l only serves as a function parameter and no sampling is done for l . The importance sampling framework that evaluates an arbitrary test function $\varphi(x)$ is given by

$$\begin{aligned} & \int \varphi(x) \frac{p_D(x, l) g(z | x, l) p_b(x, l)}{q(x | z, l)} q(x | z, l) dx \\ &= \frac{1}{N} \sum_{i=1}^N \varphi(\eta_1^i) \frac{p_D(\eta_1^i, l) g(z | \eta_1^i, l) p_b(\eta_1^i, l)}{q(\eta_1^i | z, l)}, \\ &= \frac{1}{N} \sum_{i=1}^N \varphi(\eta_1^i) \frac{p_D(\eta_1^i, l) g(z | \eta_1^i, l) p_b(\eta_1^i, l)}{p_b(\eta_0^i, l) |\det(\dot{T}^i(\eta_0^i))|}, \end{aligned} \quad (7)$$

where the last equality follows from the mapping property of the invertible flow of (5) and where the flow prior is the birth density, i.e., $p_b(\eta_0^i, l)$. From (7), we can approximate $p_b(x, l | z)$ with the set of weighted particles $\{(w^i, \eta_1^i)\}_{i=1}^N$,

where the normalized weights ($\sum_{i=1}^N w^i = 1$) are proportional to

$$w^i \propto \frac{p_D(\eta_1^i, l)g(z|\eta_1^i, l)p_b(\eta_1^i, l)}{p_b(\eta_0^i, l)}|\det(\dot{T}^i(\eta_0^i))|. \quad (8)$$

3.3. δ -GLMB persistent flow

Consider a persistent target with posterior $p_{k-1|k-1}(x, l)$ at time $k-1$. We will consider updating the target with a specific measurement $z_k \in Z_k$, and hence the superscript θ referring to the specific mapping will be dropped.

The particles $\eta_0^i \sim f_{k|k-1}(\cdot|x_{k-1}^i, l)$ are transformed via the flow into the set $\{\eta_1^i\}_{i=1}^N$. The new set of particles is effectively drawn from an importance distribution $q(\cdot|x_{k-1}^i, l, z_k)$.

According to [2, Eqs. 31 and 40], the updated density $p_{k|k}(x, l)$ is approximated by the particle set $\{(w_k^i, \eta_1^i)\}_{i=1}^N$ with associated weights w_k^i and

$$w_k^i \propto \frac{p_D(\eta_1^i, l)g(z_k|\eta_1^i, l)f_{k|k-1}(\eta_1^i|x_{k-1}^i, l)p_S(x_{k-1}^i, l)}{q(\eta_1^i|x_{k-1}^i, l, z_k)}w_{k-1}^i. \quad (9)$$

Furthermore, by employing the mapping property of the invertible flow (5), the weights become

$$w_k^i \propto \frac{p_D(\eta_1^i, l)g(z_k|\eta_1^i, l)f_{k|k-1}(\eta_1^i|x_{k-1}^i, l)p_S(x_{k-1}^i, l)}{f_{k|k-1}(\eta_0^i|x_{k-1}^i, l)} \times |\det(\dot{T}^i(\eta_0^i))|w_{k-1}^i, \quad (10)$$

where we employed the transition kernel as the prior for the flow, i.e., $p(\eta_0^i|x_{k-1}^i) = f_{k|k-1}(\eta_0^i|x_{k-1}^i)$.

3.4. Algorithm description

The δ -GLMB prediction and update steps involve an increasing number of terms in their respective summations, i.e., number of hypothesis. In practice, several truncation methodologies that avoid the explicit enumeration of all hypothesis were proposed [2, Sec. IV-V].

Prediction in the particle flow δ -GLMB is carried out in the same manner as in the standard SMC δ -GLMB prediction of [2, Sec. V], where samples are drawn from the prior birth density $p_B(\cdot, l)$ or the kinematic kernel $f_{k-1|k}(\cdot, l)$ in accordance with the type of the target, i.e., birth or persistent.

For each predicted hypothesis (I, ξ) , the update step considers first the ranked assignment problem which requires the evaluation of all associations $\theta : I \rightarrow \{0, 1, \dots, |Z_k|\}$ between the label set I and the current measurement set Z_k . The costs of these associations, as seen in [2, Eq. 24], effectively require the updated posterior, which in our case is obtained via particle flow. Hence, particle flow is used to migrate the particles of the predicted targets to form new particle sets $\{\eta_1^i\}_{i=1}^N$. The importance weights of the resulting particles are computed separately for birth (8) and persistent (10) targets. Additionally, gating is performed in order to avoid

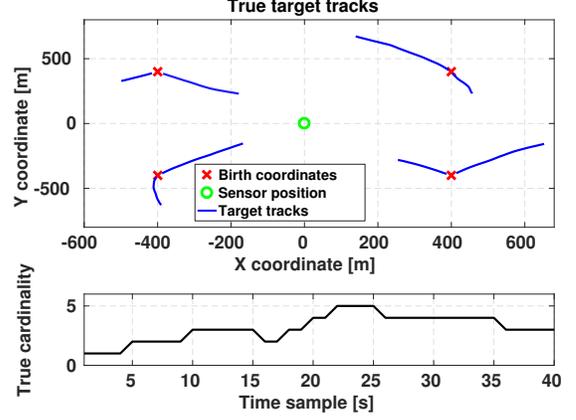


Fig. 1. The true number of targets and their tracks. Targets are born around the locations marked with \times . The sensor is marked with \circ and placed in the center of the coordinate system.

unnecessary flow operations. Regarding the flow of particles, both EDH and LEDH approaches are possible. Once the ranked assignment problem is solved, the new hypothesis can be formed as described in [2, Sec. IV.B].

4. NUMERICAL RESULTS

In this section we evaluate the performance of several implementations of the δ -GLMB filter in a non-linear tracking scenario. The δ -GLMB implementations considered are the Extended Kalman Filter (EKF), Unscented Kalman Filter (UKF) [11], the Bootstrap Particle Filter (BPF) and the Flow SMC δ -GLMB.

Targets are assumed to evolve independently of each other in a two dimensional Cartesian system. Target state vectors are taken to be $x = [p_{x,k}, p_{y,k}, \dot{p}_{x,k}, \dot{p}_{y,k}]^T$, where p_x and p_y represent the target coordinates and $\dot{p}_{x,k}$ and $\dot{p}_{y,k}$ are its velocities along the two axes. The target kinematic model induces the Gaussian transition kernel $f_{k+1|k}(x|\xi, l) = \mathcal{N}(x; \mathbf{F}_{k+1}\xi, \mathbf{Q}_{k+1})$ where \mathbf{F}_{k+1} and \mathbf{Q}_{k+1} are matrices typical of a nearly constant velocity model [12, Ch. 6.3.2]. The process noise, i.e., the acceleration noise, has a standard deviation of $\sigma = 1\text{m/s}^2$. A single simulation has a duration of 40s and the target tracks are depicted in Fig. 1. The probability of survival of targets is set to $p_S(\mathbf{x}) = 0.99$.

A single sensor is placed at coordinates $[0, 0]$ and collects bearing and range measurements, as given by

$$z_k = \left[\frac{\text{atan2}\left(\frac{p_{y,k}}{p_{x,k}}\right)}{\sqrt{(p_{x,k})^2 + (p_{y,k})^2}} \right] + w_k, \quad (11)$$

for each target where $\text{atan2}(\cdot)$ is the four-quadrant inverse tangent function and the measurement noise is $w_k \sim \mathcal{N}(0, \mathbf{R}_k)$. In the following simulation, we consider the case of $\mathbf{R} =$

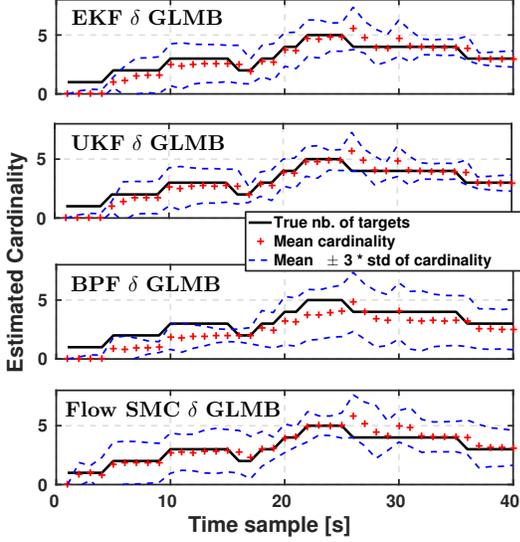


Fig. 2. Estimated cardinality of the various filters for $p_D = 0.9$.

diag(0.001, 0.001), thus insuring highly informative measurements. Target births are modeled with a labeled multi-Bernoulli [1, Eq. 12] containing 4 Bernoulli components with Gaussian probability densities located at $(\pm 400m, \pm 400m)$ and with identical $\mathbf{P} = \text{diag}(10, 10, 70, 70)$ covariance matrices. The probabilities of existence of all 4 components are set to 0.1. The BPF δ -GLMB employs 50000 particles and sampling is conducted from the kinematic kernel $f_{k+1|k}(\cdot|\cdot)$. The Flow δ -GLMB employs 1000 particles with the EDH flow variant described in Section 3. A total of 300 simulations were conducted for the tracks depicted in Fig. 1. A first result is presented in Fig. 2 for the estimated cardinality, i.e., the number of targets, for each filter. The probability

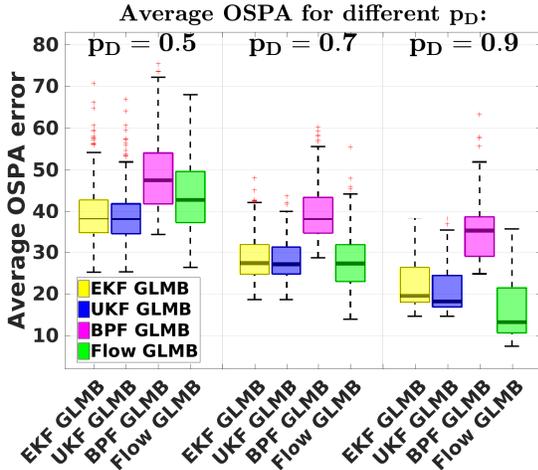


Fig. 3. OSPA error box plot at various probabilities of detection.

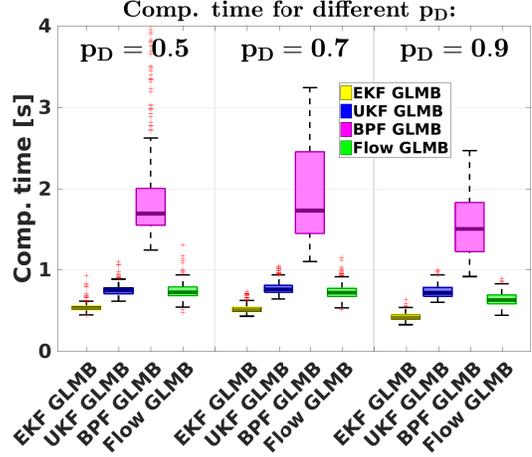


Fig. 4. Box plot for the computational time of the various filters.

of detection p_D for the sensor is set to 0.9. First notice that the BPF δ -GLMB struggles to estimate the number of targets due to the degeneracy phenomenon. The performances of the EKF and UKF δ -GLMB are similar but have difficulties with the initialization of targets. The Flow δ -GLMB does not exhibit difficulties with target initialization and has improved cardinality estimates as opposed to the BPF δ -GLMB. The Optimal Sub-Pattern Assignment (OSPA) [13] is an error metric that incorporates both cardinality mismatches and position errors between the estimated and true sets of targets. A box plot for the average OSPA errors is depicted in Fig. 3 for different p_D . The average OSPA error is defined as the mean OSPA error for a single simulation of 40s. Observe the high error reported for the BPF δ -GLMB which struggles with the highly informative measurements. Note the increased performance of the Flow δ -GLMB as opposed to BPF and also EKF/UKF δ -GLMB for high p_D . Additionally, note that the Flow δ -GLMB outperforms the BPF for all p_D . However, at low p_D the EKF and UKF implementations are more robust and outperform the particle-based methods.

A similar box-plot is given in Fig. 4 that concerns the computation time of each filter. The highest computation load is reported for the BPF δ -GLMB while the Flow, UKF and EKF implementations have relatively similar computation loads.

5. CONCLUSIONS

In this paper a particle flow SMC implementation is given for the δ -GLMB filter for multi-target tracking. The resulting Flow δ -GLMB filter offers better performance than the standard Bootstrap implementation with a computational load similar to an EKF/UKF implementation.

6. REFERENCES

- [1] Ba-Tuong Vo and Ba-Ngu Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Trans. Signal Process.*, vol. 61, no. 13, pp. 3460–3475, July 2013.
- [2] Ba-Ngu Vo, Ba-Tuong Vo, and Dinh Phung, "Labeled random finite sets and the Bayes multi-target tracking filter," *IEEE Trans. Signal Process.*, vol. 62, no. 24, pp. 6554–6567, Dec 2014.
- [3] P. Bickel, B. Li, and T. Bengtsson, "Sharp failure rates for the bootstrap particle filter in high dimensions," in *Pushing the limits of contemporary statistics: Contributions in honor of Jayanta K. Ghosh*, pp. 318–329. Institute of Mathematical Statistics, 2008.
- [4] F. Daum and J. Huang, "Nonlinear filters with with log-homotopy," in *Proc. SPIE Signal and Data Processing of Small Targets*, San Diego, CA, Sept. 2007.
- [5] P. Bunch and S. Godsill, "Approximations of the optimal importance density using Gaussian particle flow importance sampling," *J. Amer. Statist. Assoc.*, vol. 111, no. 514, pp. 748–762, 2016.
- [6] Y. Li and M. Coates, "Particle filtering with invertible particle flow," *submitted to IEEE Trans. Sig. Proc.*, arXiv:1607.08799 [stat.ME].
- [7] R.P.S. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *EEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152 – 1178, October 2003.
- [8] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*, Artech House, Norwood, MA, 2007.
- [9] F. Daum, J. Huang, and A. Noushin, "Exact particle flow for nonlinear filters," in *Proc. SPIE Conf. Signal Proc., Sensor Fusion, Target Recog.*, Orlando, FL, Apr. 2010, p. 769704.
- [10] T. Ding and M. J. Coates, "Implementation of the Daum-Huang exact-flow particle filter," in *Proc. IEEE Statistical Signal Processing Workshop (SSP)*, Ann Arbor, MI, Aug. 2012, pp. 257–260.
- [11] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proc. IEEE*, vol. 92, no. 3, pp. 401–422, March 2004.
- [12] Y. Bar-Shalom, X. Rong Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Algorithms and Software for Information Extraction.*, John Wiley and Sons, Hoboken, NJ, July 2001.
- [13] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3447–3457, August 2008.