

Optimizing Information Credibility in Social Swarming Applications

Bin Liu*, Peter Terlecky[‡], Amotz Bar-Noy[‡], Ramesh Govindan*, Michael J. Neely*

*University of Southern California, Los Angeles [‡]City University of New York

Abstract—With the advent of smartphone technology, it has become possible to conceive of entirely new classes of applications. *Social swarming*, in which users armed with smartphones are directed by a central *director* to report on events in the physical world, has several real-world applications: search and rescue, coordinated fire-fighting, and the DARPA balloon hunt challenge. In this paper, we focus on the following problem: how does the director optimize the selection of reporters to deliver credible corroborating information about an event. We first propose a model, based on common intuitions of believability, about the credibility of information. We then cast the problem posed above as a discrete optimization problem, and introduce optimal centralized solutions and an approximate solution amenable to decentralized implementation whose performance is about 20% off on average from the optimal (on real-world datasets derived from Google News) while being 3 orders of magnitude more computationally efficient. More interesting, a time-averaged version of the problem is amenable to a novel stochastic utility optimization formulation, and can be solved optimally, while in some cases yielding decentralized solutions. To our knowledge, we are the first to propose and explore the problem of extracting credible information from a network of smartphones.

I. INTRODUCTION

With the advent of smartphone technology, it has become possible to conceive of entirely new classes of applications. Recent research has considered personal reflection [1], social sensing [2], lifestyle and activity detection [3], and advanced speech and image processing applications [4]. These applications are enabled by the programmability of smartphones, their considerable computing power, and the presence of a variety of sensors on-board.

In this paper, we consider a complementary class of potential applications, enabled by the same capabilities, that we call *social swarming*. In this class of applications, a *swarm* of users, each armed with a smart phone, cooperatively and collaboratively engages in one or more tasks. These users often receive instructions from or send *reports* (a video clip, an audio report, a text message, or etc.) to a *swarm director*. Because directors have a global view of information from different users, directors are able to manage the swarm efficiently to achieve the task's objectives. Beyond the obvious military applications,

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

there are several civilian ones: search and rescue, coordinated fire-fighting, and the DARPA balloon hunt challenge¹.

In these applications, an important challenge is to obtain *credible* (or believable) information. In general, sociologists have observed three ways in which believable information might be obtained [5]: *homophily*, by which people believe like-minded people; *test-and-validate*, by which the recipient of information tests the correctness of the information; and *corroboration*, where the belief in information is reinforced by several sources reporting the same (or similar) information. The process by which humans believe information is exceedingly complex, and an extended discussion is beyond the scope of this paper.

Instead, our focus is on simple and tractable models for corroboration in social swarming type applications. Specifically, the scenario we consider is the following. Suppose that an *event* (say, a balloon sighting) is reported to a swarm director. The director would like to corroborate this report by obtaining reports from other swarm members: which reporters should she select? We call this the *corroboration pull* problem. Clearly, asking every swarm member to report is unnecessary, at best: swarms can have several hundred participants, and a video report from each of them can overwhelm the network. Thus, intuitively, the director would like to *selectively* request reports from a subset of swarm members, while managing the network resources utilized.

In this paper, we formalize this intuition and study the space of corroboration pull formulations. Our contributions are three-fold. 1) We introduce a model for the credibility of reports. This model quantifies common intuitions about the believability of information: for example, that video is more believable than text, and that a reporter closer to an event is more believable than one further away (Section II). 2) We then cast the *one-shot* corroboration pull problem as a discrete optimization problem and show that it reduces to a multiple-choice knapsack problem with weakly-polynomial optimal solutions. We develop strongly-polynomial, but inefficient, solutions for the case when the number of formats is fixed, and an optimal algorithm for the case of two formats. Finally, we derive an approximation algorithm for the general case that leverages the structure of our credibility model. This algorithm is about 20% off the optimal, but its running time is 2-3 orders of magnitude faster than the optimal algorithm, a running time difference that can make the difference between winning and losing in, say, a balloon hunt. 3) We then show

¹<http://www.crn.com/networking/222000334>

that, interestingly, the *renewals* version of the problem, where the goal is to optimize corroboration pull in a time-averaged sense, can be solved optimally, while admitting a completely decentralized solution.

II. TERMINOLOGY, MODEL, AND OPTIMIZATION FORMULATION

As smart phones proliferate, social swarming applications are likely to become increasingly common. In this paper, we consider a constrained form of a social swarming application in which N participants, whom we call *reporters*, collaboratively engage in a well-defined task. Each reporter is equipped with a smart phone and directly reports to a swarm director using the 3G/EDGE network. A reporter may either be a human being or a sensor (static, such as a fixed camera, or mobile, as a robot). A *director* (either a human being, or analytic software) assimilates these reports, and may perform some actions based on the content of these combined reports.

Our setting is simplified in many ways. For now, we consider a situation where reporters cooperate, and are therefore benign: we leave a consideration of malicious reporting to future work. Similarly, we have implicitly assumed an always available 3G/EDGE network, and have not considered network dynamics (such as the availability of opportunistic WiFi networks). We believe this assumption can be relaxed using techniques from our prior work [6], but have left an exploration of this to future work. Despite these simplifications, we show that the problem space has sufficient richness in and of itself.

Each reporter reports on an *event*. The nature of the event depends upon the social swarming application: for example, in a search and rescue operation, an event corresponds to the sighting of an individual who needs to be rescued; in the balloon hunt, an event is the sighting of a balloon. Events occur at a particular *location*, and multiple events may occur concurrently either at the same location or at different locations.

Reporters can transmit reports of an event using one of several formats: such as a video clip, an audio clip, or a text message describing what the report sees. Each report is a form of *evidence* for the existence of the event. As we discuss below, different forms of evidence are “believed” to different extents. In general, we assume that each reporter is capable of generating R different report formats, denoted by f_j , for $1 \leq j \leq R$. However, different formats have different costs to the network: for example, video or audio could consume significantly higher transmission resources than, say, text. We denote by e_j the cost of a report f_j ; for ease of exposition, we assume that reports are a fixed size so that all reports of a certain format have the same cost (our results can be easily generalized to the case where report costs are proportional to their length).

Finally, reporters can be mobile, but we assume that the director is aware of the location of each reporter. In our problem formulation, we ignore the cost of sending periodic location updates to the director. In practice, this may be a reasonable assumption for three reasons. First, the cost of location updates may be amortized over other context aware applications that may be executing on the smart phone.

Second, although this cost may be significant, it adds a fixed cost to our formulations and does not affect the results we present in the paper. Finally, the absolute cost of the location updates themselves is significantly less than the cost of video transmissions, for example.

Now, suppose that the director in a swarming application has heard, through out of band channels or from a single reporter, of the existence of an event E at location L . To verify this report, the director would like to request *corroborating* reports from other reporters in the vicinity of L . Which reporters should she get corroborating reports from? What formats should those reporters use?

To understand this, recall that the goal of corroboration is to increase the director’s belief in the occurrence of the event. How much should the director believe a specific reporter? Or, equivalently, what is the *credibility* of a report?

In general, this is a complex sociological and psychological question which, at the moment, is not objectively quantifiable. However, in this paper, we model the credibility of the report using two common intuitions about credibility. The first intuition is based on the maxim “seeing is believing”: a video report is more credible than a text report. We extend this maxim in our model to incorporate other formats, like audio: audio is generally less credible than video (because, while it gives some context about an event, video contains more context), but more credible than text (for a similar reason). Of course, this is an assumption: video and audio can be just as easily doctored as text. Recall that our model, for now, assumes cooperative non-malicious elements: in future work, we plan to discuss how to model the credibility of reports in the presence of malicious elements. Moreover, as we shall show later, many of our results are insensitive to the exact choice of the credibility model.

Our second intuition is based on the often heard statement “I’ll believe someone who was there”, suggesting that proximity of the reporter to an event increases the credibility of the report. More precisely, a report A generated by a reporter at distance d_a from an event has a higher credibility than a report B generated by a reporter at a distance d_b , if $d_a < d_b$. This is also a simplified model: the real world is more complex, since the complexity of the terrain, or line of sight, may matter more than geometric distance.

While are many different ways in which we can objectively quantify the credibility of a report given these intuitions, we picked the following formulation. Let S_i be the position of reporter i , L be the position of event E and $c_{i,j}(S_i, L)$ be the credibility of the report generated by reporter i when report format f_j is used. We define $c_{i,j}(S_i, L)$ as:

$$c_{i,j}(S_i, L) = \begin{cases} \gamma_j / d(S_i, L)^{\delta_j}, & \text{if } h_0 < d(S_i, L) \\ \gamma_j / h_0^{\delta_j}, & \text{if } d(S_i, L) \leq h_0 \end{cases} \quad (1)$$

with $1 \leq j \leq R$, $\gamma_1 \leq \gamma_2 \leq \dots \leq \gamma_R$, and $\delta_1 > \delta_2 > \dots > \delta_R$. Here, $d(\cdot)$ is the Euclidean distance between points, h_0 is a certain minimum distance to avoid division by zero as well as to bound the maximum credibility to a certain level, γ_j is a constant of proportionality implying the maximum achievable credibility of report format f_j , and the credibility decays

according to a power-law with exponent δ_j when format f_j is used.

Our credibility model incorporates the two intuitions described above as follows. The intuition about the credibility being dependent on proximity is captured by the power-law decay with distance. The intuition about the higher credibility of the video compared to text is captured by having a larger γ and a smaller exponent for video.

This model can be extended to incorporate *noise* or *confusion*. For example, poor visibility or audible noise near a reporter may, depending upon the format used, reduce the believability of a report. The intensity of point sources of noise can be modeled as a function that decays with distance:

$$G_1(S_i, O_1) = \frac{1}{[1 + d(S_i, O_1)]^{1/\sigma_1}} \quad (2)$$

where S_i is the position of reporter i , O_1 is the position of noise source 1, and σ_1 represents the strength and effective range of noise source 1. Then, if for reporter i and event E , the original credibility without noise is $c_{i,j}(S_i, L)$, then the credibility with X noise sources should be

$$c'_{i,j}(S_i, L) = c_{i,j}(S_i, L) \prod_{p=1}^X (1 - G_p(S_i, O_p)) \quad (3)$$

Noise sources effectively increase the distance of the reporter from the event, reducing his or her credibility. As we show later, our solutions can incorporate this form of noise without any modification.

Although we have assigned objective quantitative values to credibility, belief or disbelief is often qualitative and subjective. Thus, we don't expect swarm directors to make decisions based on the exact values of credibility of different reports, but rather to operate in one of two modes: *a*) ask the network to deliver corroborating reports whose total credibility is above a certain threshold, while minimizing cost, or *b*) obtain as much corroborating information that they can get from the network for a given cost. We study these two formulations, respectively called MINCOST and MAXCRED.

Before doing so, there are two questions to be answered: What is the value of the credibility of a collection of corroborating reports? What is the physical/intuitive meaning of a threshold on the credibility? For the first question, there are many possible answers and we consider two. With an *additive corroboration function*, the total credibility is simply the sum of the individual credibilities. More generally, with a *monotonically-increasing corroboration function*, the total credibility increases monotonically as a function of the sum of the individual credibilities. The second question is important because it can help directors set thresholds appropriately. The intuition for a particular threshold value C can be explained as follows. Suppose a director would be subjectively satisfied with 3 corroborating video clips from someone within 10m of an event. One could translate this subjective specification into a threshold value by simply taking the sum of the credibilities of 3 video reports from a distance of 10m.

In the next two sections, we formally define MINCOST and MAXCRED, and then consider two problem variants: a *one-shot* problem which seeks to optimize reporting for individual

TABLE I
NOTATION

N	the total number of available reporters
$c_{i,j}$	the short form of (1) in a given event
R	the total number of report formats
e_j	the cost when using report format f_j
C	the target credibility in MINCOST
A	the dynamic programming process of MINCOST
B	the cost budget in MAXCRED
D	the dynamic programming process of MAXCRED

events, and a *renewals* problem which optimizes reporting over a sequence of event arrivals.

III. THE ONE-SHOT PROBLEM

In this section, we formally state the MINCOST and MAXCRED formulations for the additive corroboration function and in the absence of noise, discuss their complexity, develop optimal solutions for them, and then explore an approximation algorithm that leverages the structure of the credibility function for efficiency. We conclude with a discussion of extensions to the formulations for incorporating the impact of noise sources, and for monotonically-increasing corroboration function. Our exposition follows the notation developed in the previous section, and summarized in Table I.

A. MINCOST and MAXCRED: Problem Formulation and Complexity

1) *Problem Formulations*: Recall that, in Section II, we informally defined the MINCOST problem to be: what is the minimum cost that guarantees total credibility $C > 0$? MINCOST can be stated formally as an optimization problem:

$$\begin{aligned} \text{Minimize :} & \quad \sum_{i=1}^N \sum_{j=1}^R x_{i,j} e_j & (4) \\ \text{Subject to:} & \quad \sum_{i=1}^N \sum_{j=1}^R x_{i,j} c_{i,j} \geq C \\ & \quad x_{i,j} \in \{0, 1\}, \forall i \in \{1, \dots, N\}, \forall j \in \{1, \dots, R\} \\ & \quad \sum_{j=1}^R x_{i,j} \leq 1, \forall i \in \{1, \dots, N\} \end{aligned}$$

where $x_{i,j}$ is a binary variable that is 1 if reporter i uses format f_j , and 0 otherwise.

Analogously, we can formulate MAXCRED (the maximum credibility that can be achieved for a cost budget of $B > 0$) as the following optimization problem:

$$\begin{aligned} \text{Maximize :} & \quad \sum_{i=1}^N \sum_{j=1}^R x_{i,j} c_{i,j} & (5) \\ \text{Subject to:} & \quad \sum_{i=1}^N \sum_{j=1}^R x_{i,j} e_j \leq B \\ & \quad x_{i,j} \in \{0, 1\}, \forall i \in \{1, \dots, N\}, \forall j \in \{1, \dots, R\} \\ & \quad \sum_{j=1}^R x_{i,j} \leq 1, \forall i \in \{1, \dots, N\} \end{aligned}$$

2) *On the Complexity of MINCOST and MAXCRED*: If, in the above formulation, the cost e_j is also dependent on the identity of the reporter (and therefore denoted by $e_{i,j}$), the MAXCRED problem can be shown to be a special instance of the Multiple-Choice Knapsack Problem (MCKP, [7]). Moreover, the special case of one format (and $e_{i,j} = e_i$) is the well-known Knapsack problem (KP) which is NP-hard. However, when the cost is dependent only on the format (i.e., $e_{i,j} = e_j$), we can state the following theorem, whose proof (omitted for brevity) uses a reduction from the original Knapsack problem.

Theorem 3.1: MINCOST and MAXCRED are NP-Hard.

B. Optimal Solutions

Despite Theorem 3.1, it is instructive to consider optimal solutions for the two problems for two reasons. First, for many social swarming problem instances, the problem sizes may be small enough that optimal solutions might apply. Second, optimal solutions can be used to calibrate an approximation algorithm that we discuss later. In this section, we discuss two classes of optimal solutions for MINCOST and MAXCRED, with different tradeoffs: one based on *dynamic programming*, and another based on a *min-cost flow formulation*.

1) *Dynamic Programming*: Since there exist optimal, weakly-polynomial algorithms for MCKP, it is natural that similar algorithms exist for MINCOST and MAXCRED. We describe these algorithms for completeness, since we use them in a later evaluation.

For MINCOST (4), we can write $y_{i,j} = 1 - x_{i,j}$, where $y_{i,j} \in \{0, 1\}$, and then we have:

$$\sum_{i=1}^N \sum_{j=1}^R e_j - \text{Maximize} \sum_{i=1}^N \sum_{j=1}^R y_{i,j} e_j \quad (6)$$

$$\begin{aligned} \text{Subject to:} \quad & \sum_{i=1}^N \sum_{j=1}^R y_{i,j} c_{i,j} \leq \sum_{i=1}^N \sum_{j=1}^R c_{i,j} - C = W \\ & y_{i,j} \in \{0, 1\}, \forall i \in \{1, \dots, N\}, \forall j \in \{1, \dots, R\} \\ & \sum_{j=1}^R y_{i,j} \geq R - 1, \forall i \in \{1, \dots, N\} \end{aligned}$$

where the minimization problem (4) has been transformed into a maximization problem, and the notation in (6) emphasizes that the first term in the total cost $\sum_{i=1}^N \sum_{j=1}^R e_j$ does not depend on the $y_{i,j}$ variables to be optimized. For a given event, the sum of the $c_{i,j}$ values is a constant, and so W is also a constant.

This optimization problem can be solved by a dynamic programming approach if we assume all $c_{i,j}$ s are truncated to a certain decimal precision, so that $c_{i,j} \in \{0, \zeta, 2\zeta, \dots\}$ where ζ is a discretization unit. Then for any binary $y_{i,j}$ values that meet the constraints of the above problem, the sum $\sum_{i=1}^N \sum_{j=1}^R y_{i,j} c_{i,j}$ takes values in a set $\mathcal{W} \triangleq \{0, \zeta, 2\zeta, \dots, W\}$. Note that the cardinality $|\mathcal{W}|$ depends on N , R , the $c_{i,j}$ values, and the discretization unit ζ . Now define $A(l, s)$ as the subproblem of selecting reporters in the set $\{1, \dots, l\}$ subject to a constraint s . Assuming $A(l, s)$ values are known for a particular l , we recursively compute $A(l+1, s)$ for all $s \in \mathcal{W}$ by:

$$A(l+1, s) = \max\{\phi^{(0)}(l, s), \phi^{(1)}(l, s), \dots, \phi^{(R)}(l, s)\} \quad (7)$$

where $\phi^{(k)}(l, s)$ is defined for $k \in \{0, 1, \dots, R\}$:

$$\phi^{(k)}(l, s) \triangleq A(l, s - \sum_{j=1, j \neq k}^R c_{l,j}) + \sum_{j=1, j \neq k}^R e_j$$

This can be understood as follows: The value $\phi^{(k)}(l, s)$ is the cost associated with reporter $l+1$ using option $k \in \{0, 1, \dots, R\}$ and then allocating reporters $\{1, \dots, l\}$ according to the optimal solution $A(l, s - \sum_{j=1, j \neq k}^R c_{l,j})$ that corresponds to a smaller budget. Note that option $k \in \{1, \dots, R\}$ corresponds to reporter $l+1$ using a particular format (so that $y_{l+1,k} = 0$ for option k and $y_{l+1,m} = 1$ for all $m \neq k$), and option $k = 0$ corresponds to reporter $l+1$ remaining idle (so that $y_{l+1,m} = 1$ for all m). The time complexity of this dynamic programming algorithm, called MINCOST-DP, is $O(NR|\mathcal{W}|)$.

Similarly, MAXCRED can be solved using dynamic programming, yielding an algorithm we label MAXCRED-DP:

$$\begin{aligned} D(l+1, s) &= \max[D(l, s), \mu^{(1)}(l, s), \mu^{(2)}(l, s), \dots, \mu^{(R)}(l, s)], \\ \mu^{(k)}(l, s) &\triangleq D(l, s - e_k) + c_{l,k} \text{ for } k \in \{1, \dots, R\} \end{aligned} \quad (8)$$

with complexity $O(NR|\mathcal{B}|)$. $\mathcal{B} \triangleq \{0, \eta, 2\eta, \dots, B\}$, where η is a discretization unit.

2) *Min-Cost Flow*: For a fixed number of formats, it is possible to define strongly-polynomial, but not necessarily efficient, optimal algorithms for MINCOST and MAXCRED. These solutions are based on minimum-cost flow algorithms [8]. Define α_j to be the number of reporters reporting with format f_j . Define a report vector to be $(\alpha_1, \alpha_2, \dots, \alpha_R)$ and an $(\alpha_1, \alpha_2, \dots, \alpha_R)$ -assignment to be an assignment of formats to reporters with α_j reporters reporting with format f_j for each $j \in \{1, \dots, R\}$. We shall find an $(\alpha_1, \alpha_2, \dots, \alpha_R)$ -assignment of formats to reporters of maximum credibility.

We shall do so by transforming this problem to a min-cost flow problem and applying a min-cost flow algorithm to obtain the assignment of maximum credibility, in the following manner. Assign nodes for each of the N reporters and each of the R formats. Form a complete bipartite graph between the reporters and formats. Assign the edge connecting reporter i to format f_j with max capacity 1, min capacity 0, and cost $\max_{k \in \{1, \dots, N\}} \max_{l \in \{1, \dots, R\}} c_{k,l} - c_{i,j}$. Note that minimizing the set of such costs maximizes the set of $\{c_{i,j}\}$. Also, create a source node and a sink node. Connect the source node to each of the reporter nodes and give each edge max capacity 1, min capacity 0, and cost 0. Connect the sink node to each of the format nodes and give the edge connecting to format f_j max capacity α_j , min capacity α_j , and cost 0. We shall call this network the *credibility network*. This network has $O(N+R)$ vertices and $O(NR)$ edges.

This construction ensures that applying a min-cost flow algorithm to the credibility network gives a minimum cost and maximum credibility $(\alpha_1, \dots, \alpha_R)$ -assignment of formats to reporters. Using this construction, it is fairly easy to define an optimal algorithm for MAXCRED, that we call MAXCRED-MCF. In this algorithm, we simply enumerate all possible $(\alpha_1, \dots, \alpha_R)$ -assignments, and find the highest total credibility assignment that satisfies the cost budget B . In a similar way, one can define MINCOST-MCF.

There are $O(N^{R-1})$ possible report vectors. The Enhanced Capacity Scaling algorithm [8] solves the minimum cost flow

problem in time $O(|E| \log |V| (|E| + |V| \log |V|))$. Thus, the Enhanced Capacity Scaling algorithm runs in time $O(NR \log(N+R)(NR + ((N+R) \log(N+R))))$ on the credibility network. This leads to the following lemma:

Lemma 3.2: Both MINCOST-MCF and MAXCRED-MCF run in time $O(N^R R \log(N+R)(NR + ((N+R) \log(N+R))))$.

Note that when the number of formats R is fixed, these algorithms are polynomial in N . In addition, when $|\mathcal{B}|, |\mathcal{W}| = \omega(N^{R-1} \log(N+R)(NR + ((N+R) \log(N+R))))$ these algorithms have lower asymptotic complexity than their dynamic programming equivalents.

C. Leveraging the Structure of the Credibility Function

The solutions discussed so far do not leverage any *structure* in the problem. Given an event and reporter locations, the credibility associated with each report format is computed as a number and acts as an input to the algorithms discussed. However, there are two interesting structural properties in the problem formulation. First, for a given reporter at a given location, the credibility is higher for a format whose cost is also higher. Second, for reporters at different distances, the credibility decays as a function of distance. In this section, we ask the question: can we leverage this structure to devise efficient approximation algorithms, or optimal special-case solutions either for MAXCRED or MINCOST?

1) *An Efficient Optimal Greedy MAXCRED Algorithm for Two Formats:* When a social swarming application only uses two report formats (say, text and video), it is possible to devise an optimal greedy MAXCRED algorithm. Assume each of the N reporters can report with one of two formats, f_1 or f_2 , that reporters are indexed so that reporter i is closer to the event than reporter k , for $i < k$, and that credibility decays with distance. Furthermore, we assume that $e_1 = \beta > 1$ and $e_2 = 1$ and that $c_{i,1} \geq c_{i,2} \forall i \in \{1, \dots, N\}$: that the more expensive format yields a higher credibility.

With these assumptions, the following algorithm, denoted MAXCRED-2F, finds an assignment with maximum credibility that falls within a budget B and runs in time $O(N^2)$.

Algorithm 1 Algorithm MAXCRED-2F

INPUT: $(c_{i,j})$: $i \in \{1, \dots, N\}, j \in \{1, 2\}$; $(1, \beta)$; Budget B

Define $d_m \triangleq c_{m,1} - c_{m,2}$ for each $m \in \{1, \dots, N\}$.

For $i \in \{0, \dots, \min[\lfloor B/\beta \rfloor, N]\}$, do:

- 1) Define $Y \triangleq \min[N - i, \lfloor B - \beta i \rfloor]$.
- 2) Define the *active set* $\mathcal{A} \triangleq \{1, \dots, i + Y\}$, being the set of $i + Y$ reporters closest to the event.
- 3) Define \mathcal{D}^* as the set of i reporters in \mathcal{A} with the largest d_m values (breaking ties arbitrarily). Then choose format f_1 for all reporters $m \in \mathcal{D}^*$, choose f_2 for $m \in \mathcal{A} - \mathcal{D}^*$, and choose “idle” for all $m \notin \mathcal{A}$.
- 4) Define C_{MAX}^i as the total credibility of this assignment:

$$C_{MAX}^i \triangleq \sum_{m \in \mathcal{A}} c_{m,2} + \sum_{m \in \mathcal{D}^*} d_m$$

OUTPUT: $i^* \triangleq \arg \max_i C_{MAX}^i$.

The output of this algorithm is the maximum credibility assignment of formats to reporters. We can prove that this algorithm is optimal.

Theorem 1: The above algorithm finds the solution C_{MAX} to MAXCRED-2F problem with budget B .

Proof: For each i , we first seek to find C_{MAX}^i , the maximum credibility subject to having exactly i reporters use the expensive format f_1 . Using a simple interchange argument together with the fact that credibility of each format is non-negative and non-increasing in distance, we can show that there exists an optimal solution that activates the set \mathcal{A} that consists of $i + Y$ reporters closest to the event. Indeed, if an optimal solution does not use the set \mathcal{A} , we can swap an idle reporter closer to the event with an active reporter further from the event, without affecting cost or decreasing credibility.

For each subset $\mathcal{D} \subseteq \mathcal{A}$ that contains i reporters, define $C(\mathcal{D})$ as the credibility of the assignment that assigns reporters $m \in \mathcal{D}$ the format f_1 , assigns the remaining reporters in \mathcal{A} the format f_2 , and keeps all reporters $m \notin \mathcal{A}$ idle:

$$C(\mathcal{D}) = \sum_{m \in \mathcal{A}} c_{m,2} + \sum_{m \in \mathcal{D}} d_m$$

Then $C(\mathcal{D})$ is maximized by the subset \mathcal{D}^* containing the i reporters in \mathcal{A} with the largest d_m values. This defines C_{MAX}^i , and C_{MAX} is found by maximizing over all possible i . ■

We can analogously define a MINCOST version for two formats, but omit it for brevity. Currently, we have not been able to extend this type of algorithm to 3 formats and beyond, so this remains an interesting open problem.

2) *A Computationally-Efficient Approximation Algorithm:* The structure of our credibility function can also be used to reduce computational complexity. To understand this, recall that the dynamic programming algorithms described above jointly optimized both reporter selection and format selection. In this section, we describe an approximation algorithm for MINCOST, called MINCOST-CC, where the structure of the credibility function is used to determine, for each reporter, the format that the reporter should use. As we shall show, MINCOST-CC has significantly lower run-times at the expense of slight non-optimality in its results.

MINCOST-CC is based on the following intuition. Close to the location of the event, even low-cost formats have reasonable credibility. However, beyond a certain distance, the credibility of low-cost formats like text degrades significantly, to the point where even the small cost of that format may not be justified. Put another way, it is beneficial for a reporter to use that format whose credibility per unit cost (hence MINCOST-CC) is highest — this gives the most “bang for the buck”. Thus, for a given reporter, its current distance d from the location of the event may pre-determine the format it uses. Of course, this pre-determination can result in a non-optimal choice, which is why MINCOST-CC is an approximation algorithm.

Formally, in MINCOST-CC, if, for a reporter i :

$$k^* = \arg \max_k \left[\frac{c_{i,k}(S_i, L)}{e_k} \right]$$

then reporter i chooses format f_{k^*} . This choice can be pre-computed (since it depends only upon the credibility and cost models), but each reporter needs to recalculate its choice of the report format whenever its relative distance to the concerning event changes. The event locations that determine the format

f_{k^*} chosen by a particular reporter i form annular regions about the reporter.

Once each reporter has made the format choice, it remains for the director to decide which reporter(s) to select. For MINCOST-CC, the minimum cost formulation is identical to (6), and with comparable complexity, but with two crucial differences: both the constant $|\mathcal{W}|$ and the runtime now relate only to the number N of reporters, not to $N \times R$. As we shall show below, this makes a significant practical difference in runtime, even for moderate-sized inputs.

In MINCOST-CC, the dynamic programming process of (7) is replaced by

$$A(l+1, s) = \max \{A(l, s), c_l + A(l, s - e_l)\} \quad (9)$$

where c_l replaces $c_{l,j}$ in (7), since each reporter precomputes its format of choice. Compared with (7), the time complexity of (9) is reduced to $O(N|\mathcal{W}|)$ with a much smaller $|\mathcal{W}|$ in general. Notice that this time complexity is independent of R , the number of report formats, greatly improving its computational efficiency at the expense of some optimality.

Using steps similar to that presented in Section III-B, it is possible to define a MAXCRED-CC approximation algorithm for maximizing credibility. We omit the details for brevity, but indicate that MINCOST-CC and MAXCRED-CC still have weakly-polynomial asymptotic complexity, but are computationally much more efficient than MINCOST-DP and MAXCRED-DP.

Evaluation of MINCOST-CC. The approximation algorithms discussed above trade-off optimality for reduced computational complexity. As such, it is important to quantify this trade-off for practical swarm configurations. In this section, we compare MINCOST-CC with MINCOST-DP²: lack of space prevents us from a more extensive evaluation, but we expect our conclusions to hold in general.

Lacking data from social swarming applications, we use two different data sets. First, we carefully³ manually mine Google News for interesting events. Searching for a specific set of keywords describing an event in Google News retrieves a list of news items related to that event within 24 hours of occurrence of that event. The event location is explicitly specified in the news items. Each news item has a location, which is assumed to be the location of a reporter. We use the event location and report location as inputs to MINCOST-CC and MINCOST-DP. In this paper, we present results from three events: an event of *regional* scope, a basketball playoff game between the Lakers and the Jazz (31 reporters); an event of *national* scope, the passage of the healthcare reform bill (63 reporters); and an event of *global* scope, the opening of the Shanghai exposition (88 reporters). Of course, this choice of a surrogate for social swarming is far from perfect. However, this data set gives a varied, realistic reporter location distribution; since our algorithms depend heavily on location, we can draw some reasonable conclusions about their relative performance. That said, we also use a dataset generated from

²In some of our evaluations, we use $R = 4$. In this regime, MINCOST-DP is more efficient than MINCOST-MCF, hence the choice.

³We re-scaled reporter distances and did several data cleaning operations: removing blog posts, handling duplicate reports etc. We omit a discussion of these for brevity.




Fig. 1. Minimal cost of 4 formats with increasing k

a random distribution of reporters to ensure that we are not misled by the Google News data set, but also to explore the impact of larger swarm sizes.

We are interested in two metrics: the *optimality gap*, which is the ratio of the min-cost obtained by MINCOST-CC to that obtained by MINCOST-DP; and the *runtime* of the computation for each of these algorithms.

Figure 1 plots these two metrics as a function of the credibility threshold, expressed as a number k . A value k represents a credibility threshold corresponding to the total credibility of k reports of the highest cost format from a distance h_0 (e.g., if k is 3 and the highest cost format is video, then the director is interested in obtaining credibility equivalent to that from three video reports). In this graph, we use four data formats with $h_0, \gamma_{1-4}, \delta_{1-4}$ and the corresponding e_{1-4} setting to 1, (1, 1, 1, 1), (2, 1.5, 1, 0.5) and (1, 2.2, 5.4, 13.7) respectively. We have evaluated different numbers of data formats and different parameter settings and have obtained qualitatively similar results, but omit a discussion of these for lack of space.

From Figure 1(a), the optimality gap is, on average 19.7%, across different values of k . This is encouraging, since it suggests that MINCOST-CC produces results that are not significantly far from the optimal. Interestingly, no optimal solution exists for $k > 5$ for the regional event: this credibility threshold experiences a “saturation”, since there does not exist a set of reporters who can collectively satisfy that threshold. Other events saturate at different values of k . Finally, while this is not apparent from these graphs, the minimum cost solution is approximately linear in k for MINCOST-CC and MINCOST-DP.

More interestingly, from Figure 1(b), it is clear that the runtime of MINCOST-CC is 2-3 orders of magnitude lower than that of MINCOST-DP with the discretization setting $|\mathcal{W}| = 1000W$. This difference is not just a matter of degree, but may make the difference between a useful application and one that is not useful: MINCOST-DP can take several *tens of seconds* to complete while MINCOST-CC takes at most a few hundred milliseconds, which might make the difference between victory and defeat in a balloon hunt, or life and death in a disaster response swarm! The explanation for the performance difference is the lower asymptotic complexity of MINCOST-CC. A subtle finding is that the running time of both MINCOST-CC and MINCOST-DP decreases, sometimes dramatically in the case of MINCOST-CC, with increasing k . Intuitively, this is because there are fewer candidate sets of reporters who can satisfy a higher credibility, resulting in a smaller search space.




Fig. 2. Minimal cost in random topologies with increasing k (error bars are very small thus ignored in (b))

For random topologies, Figure 2 plots the optimality gap and runtime, averaged over 50 simulations. MINCOST-CC is, on average 20.5% and 17.4% for 100 and 200 reporters, off the optimal for different values of k , but is still 2-3 orders of magnitude more efficient than MINCOST-DP . The runtimes for both algorithms are slightly higher, given the larger number of reporters. Moreover, with 100 or 200 reporters, the optimality gap has the same upper bound, about 35% for large k . This is also observed in other simulations for report numbers of 50, 150 and 300 (not shown). We have left an analytical exploration of this upper bound to future work. Finally, a comparison of these results with Figure 1(a) reveals an interesting result. Although different types of reporter deployments can result in different optimality gap curves (the curves for the three different types of Google News in Figure 1(a) are not the same), the national event seems a qualitatively similar optimality gap curve as the random topologies, suggesting that its deployment is similar to that event. Understanding this in greater depth is also left to future work.

D. Extensions

Incorporating sources of noise into our algorithms is straightforward, so we will mention this briefly. Recall that the way we model a noise source increases a reporter's effective distance. Since our optimal algorithms, like MINCOST-DP or MINCOST-MCF , are agnostic to the structure of the credibility function, they are unaffected by noise. For an algorithm like MINCOST-CC , which does take structure into account, recall that noise sources increase a reporter's effective distance. Since reporters can quantify ambient noise, they can each use the effective distance to calculate the report format to use.

Finally, our algorithms can, in general, deal with monotonically increasing corroboration functions where the total credibility of a collection of reporters may be a non-linear function of the individual credibilities. If $I(\cdot)$ were to represent a monotonically increasing credibility function, we only need use $I(c)$ to replace c in our dynamic programming formulation. For example, (9) would become

$$A(l+1, s) = \max \{A(l, s), c_l + A(l, I(s - e_l))\}$$

Similar changes can be applied to other dynamic programming formulations.

IV. THE RENEWALS PROBLEM: RANDOMLY ARRIVING EVENTS

In the previous section, we discussed a one-shot problem: that of optimizing a single event. We now consider a sequence

of events that arrive at times $\{t_1, t_2, t_3, \dots\}$, where t_k is a real number that represents the arrival time of event k . We assume that $t_k < t_{k+1}$ for all k . In this setting, we consider a stochastic variant of MAXCRED , called $\text{MAXCRED-STOCHASTIC}$: Instead of maximizing credibility for a single event subject to a cost constraint, we maximize the average credibility-per-event subject to an average cost constraint and a per-event credibility minimum. This couples the decisions needed for each event. However, we first show that this time average problem can be solved by a reduction to individual knapsack problems of the type described in previous sections. We then show that if the per-event credibility minimum is removed, then *decisions can be made in a decentralized fashion*. Specifically, after the processing of every event, the swarm director passes a weight to all reporters. The reporters then make uncoordinated decisions when processing the next event, without any intervention from the swarm director. We derive a similar distributed version for stochastic MINCOST , labelled $\text{MINCOST-STOCHASTIC}$.

We start by solving the general time average problem using Lyapunov optimization [9], which can handle $\text{MAXCRED-STOCHASTIC}$, $\text{MINCOST-STOCHASTIC}$, as well as variations with more general constraints.

A. The General Stochastic Problem

Let $\omega[k]$ represent a random vector of parameters associated with each event k , such as the location of the event and the corresponding costs and credibilities. While $\omega[k]$ can include different parameters for different types of problems, we shall soon use $\omega[k] \triangleq [(c_{i,j}[k]), (e_j[k])]$, where $(c_{i,j}[k])$ is the matrix of event- k credibility values for reporters $i \in \{1, \dots, N\}$ and formats $f_j \in \{f_1, \dots, f_R\}$, and $(e_j[k])$ is a vector of cost information. We assume the process $\omega[k]$ is *ergodic* with a well defined steady-state distribution. The simplest example is when $\omega[k]$ is independent and identically distributed (i.i.d.) over events $k \in \{1, 2, 3, \dots\}$.

Let frame k denote the period of time $[t_k, t_{k+1})$ which starts with the arrival of event k and ends just before the next event. For every frame k , the director observes $\omega[k]$ and chooses a *control action* $\alpha[k]$ from a general set of feasible actions $\mathcal{A}_{\omega[k]}$ that possibly depend on $\omega[k]$. The values $\omega[k]$ and $\alpha[k]$ together determine an $M + 1$ dimensional vector $\mathbf{y}[k]$, representing *network attributes* for event k :

$$\mathbf{y}[k] = (y_0[k], y_1[k], \dots, y_M[k])$$

Specifically, each $y_m[k]$ attribute is given by a general function of $\alpha[k]$ and $\omega[k]$:

$$y_m[k] = \hat{y}_m(\alpha[k], \omega[k]) \quad \forall m \in \{0, 1, \dots, M\}$$

The functions $\hat{y}_m(\alpha[k], \omega[k])$ are arbitrary (possibly non-linear, non-convex, discontinuous), and are only assumed to be bounded.

Define \bar{y}_m as the time average expectation of the attribute $y_m[k]$, averaged over all frames (assuming temporarily that the limit exists):

$$\bar{y}_m \triangleq \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{E} \{y_m[k]\}$$

The general problem is to find an algorithm for choosing control actions $\alpha[k]$ for each frame $k \in \{1, 2, 3, \dots\}$ to solve:

$$\text{Minimize: } \quad \bar{y}_0 \quad (10)$$

$$\text{Subject to: } \quad 1) \quad \bar{y}_m \leq 0 \quad \forall m \in \{1, 2, \dots, M\} \quad (11)$$

$$2) \quad \alpha[k] \in \mathcal{A}_{\omega[k]} \quad \forall \text{ frames } k \in \{1, 2, \dots\} \quad (12)$$

The solution to the general problem is given in terms of a positive parameter V , which affects a performance tradeoff. Specifically, for each of the M time average inequality constraints $\bar{y}_m \leq 0$, define a *virtual queue* $Z_m[k]$ with $Z_m[0] = 0$, and with frame-update equation:

$$Z_m[k+1] = \max[Z_m[k] + y_m[k], 0] \quad (13)$$

Then every frame k , observe the value of $\omega[k]$ and perform the following actions:

- Choose $\alpha[k] \in \mathcal{A}_{\omega[k]}$ to minimize:

$$V\hat{y}_0(\alpha[k], \omega[k]) + \sum_{m=1}^M Z_m[k]\hat{y}_m(\alpha[k], \omega[k])$$

- Update the virtual queues $Z_m[k]$ according to (13), using the values $y_m[k] = \hat{y}_m(\alpha[k], \omega[k])$ determined from the above minimization.

Assuming the problem is feasible (so that it is possible to meet the time average inequality constraints), this algorithm will *also* meet all of these constraints, and will achieve a time average value \bar{y}_0 that is within $O(1/V)$ of the optimum. Typically, the V parameter also affects the average size of the virtual queues (these can be shown to be $O(V)$, which directly affects the *convergence time* needed for the time averages to be close to their limiting values). The proofs of these claims follow the theory developed in [9]–[11], with minor notational adjustments needed to change the timeslot averages there to frame-averages here. Specifically, the work in [9]–[11] considers i.i.d. events $\omega[k]$, but the same holds for more general ergodic events [12].

B. Corroboration Pull as a Dynamic Optimization Problem

Here we formulate MAXCRED-STOCHASTIC. Define $\omega[k] \triangleq [(c_{i,j}[k]), (e_j[k])]$ (representing costs and credibilities of event k), and define $\alpha[k] \triangleq (x_{i,j}[k])$, where $x_{i,j}[k]$ is a binary variable that is 1 if reporter $i \in \{1, \dots, N\}$ uses format $f_j \in \{f_1, \dots, f_R\}$ on frame k . The goal is to maximize the average credibility-per-frame subject to average cost constraints and to a minimum credibility level required on each frame $k \in \{1, 2, \dots\}$:

$$\text{Maximize: } \quad \bar{c} \quad (14)$$

$$\text{Subject to: } \quad \bar{e} \leq e_{av} \quad (15)$$

$$\sum_{i=1}^N \sum_{j=1}^R x_{i,j}[k] c_{i,j}[k] \geq c_{min} \quad \forall \text{ frames } k \quad (16)$$

$$x_{i,j}[k] \in \{0, 1\} \quad \forall i, j, \forall \text{ frames } k \quad (17)$$

$$\sum_{j=1}^R x_{i,j}[k] \leq 1 \quad \forall j, \forall \text{ frames } k \quad (18)$$

where e_{av} and c_{min} are given non-negative constants, and \bar{c} and \bar{e} are defined:

$$\bar{c} \triangleq \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^N \sum_{j=1}^R \mathbb{E} \{x_{i,j}[k] c_{i,j}[k]\}$$

$$\bar{e} \triangleq \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^N \sum_{j=1}^R \mathbb{E} \{x_{i,j}[k] e_j[k]\}$$

This problem fits the general stochastic optimization framework of the previous subsection by defining $y_0(t), y_1(t)$ by:

$$y_0(t) = \hat{y}_0(\alpha(t), \omega(t)) \triangleq - \sum_{i=1}^N \sum_{j=1}^R x_{i,j}[k] c_{i,j}[k]$$

$$y_1(t) = \hat{y}_1(\alpha(t), \omega(t)) \triangleq - e_{av} + \sum_{i=1}^N \sum_{j=1}^R x_{i,j}[k] e_j[k]$$

and by defining the set $\mathcal{A}_{\omega[k]}$ as the set of all $(x_{i,j}[k])$ matrices that satisfy the constraints (16)–(18). The resulting stochastic algorithm thus defines a virtual queue $Z_1(t)$ with update:

$$Z_1[k+1] = \max \left[Z_1[k] - e_{av} + \sum_{i=1}^N \sum_{j=1}^R x_{i,j}[k] e_j[k], 0 \right]$$

and then observing the $\omega[k]$ parameters every frame k and choosing $(x_{i,j}[k])$ to solve:

$$\text{Minimize: } \quad \sum_{i=1}^N \sum_{j=1}^R x_{i,j}[k] [Z_1[k] e_j[k] - V c_{i,j}[k]] \quad (19)$$

$$\text{Subject to: } \quad \sum_{i=1}^N \sum_{j=1}^R x_{i,j}[k] c_{i,j}[k] \geq c_{min} \quad (20)$$

$$x_{i,j}[k] \in \{0, 1\} \quad \forall i, j, \sum_{j=1}^R x_{i,j}[k] \leq 1 \quad (21)$$

MAXCRED-STOCHASTIC problem has the exact same structure as the one-shot MINCOST problem described in previous sections, with the exception that the cost weights are changed from $e_j[k]$ to $Z_1[k] e_j[k] - V c_{i,j}[k]$, which can possibly be negative. However, the same knapsack technique can be used to solve it (possibly by shifting the function to be minimized by a positive constant to make the resulting terms non-negative). Further, we note that the performance of the stochastic algorithm degrades gracefully when *approximate* implementations are used, such as implementations that are off from the optimal knapsack problem by a multiplicative constant [9]. Thus, the simple MINCOST-CC heuristic can be used here.

A simple and exact distributed implementation arises if the c_{min} constraint (16) is removed (i.e., if $c_{min} \triangleq 0$). In this case the frame k decisions (19)–(21) are *separable over reporters* and reduce to having each reporter $i \in \{1, \dots, N\}$ solve:

$$\text{Minimize: } \quad \sum_{j=1}^R x_{i,j}[k] [Z_1[k] e_j[k] - V c_{i,j}[k]]$$

$$\text{Subject to: } \quad x_{i,j}[k] \in \{0, 1\} \quad \forall j, \sum_{j=1}^R x_{i,j}[k] \leq 1$$

That is, each reporter i chooses the single format $f_j \in \{f_1, \dots, f_R\}$ with the smallest value of $Z_1[k] e_j[k] - V c_{i,j}[k]$, breaking ties arbitrarily and choosing to be idle (with $x_{i,j}[k] = 0$ for all $j \in \{1, \dots, R\}$) if all of the weights $Z_1[k] e_j[k] - V c_{i,j}[k]$ are positive. The swarm director observes the outcomes of the decisions on frame k and iterates the $Z_1[k]$ update, passing the weight $Z_1[k+1]$ to all reporters before the next event occurs.

C. MINCOST-STOCHASTIC

MINCOST-STOCHASTIC can be formulated as follows:

$$\begin{aligned} \text{Minimize:} & \quad \bar{c} \\ \text{Subject to:} & \quad \bar{c} \geq c_{av}, \text{ and constraints (17), (18)} \end{aligned}$$

We can thus define $y_0(t)$ and $y_1(t)$ as:

$$\begin{aligned} y_0[k] &\triangleq \sum_{i=1}^N \sum_{j=1}^R x_{i,j}[k] e_j[k] \\ y_1[k] &\triangleq c_{av} - \sum_{i=1}^N \sum_{j=1}^R x_{i,j}[k] c_{i,j}[k] \end{aligned}$$

from which a similar distributed solution can be obtained.

V. RELATED WORK

We are not aware of any prior work in the wireless networking literature that has tackled information credibility assessment.

However, other fields have actively explored credibility, defined as the believability of sources or information [13], [15], [16]. Credibility has been investigated in a number of fields including information science, human communication, human-computer interaction (HCI), marketing, psychology and so on [17]. In general, research has focused on two threads: the factors that affect credibility, and the dynamics of information credibility.

The seminal work of Hovland et al. [14] may be the earliest attempt on exploring credibility, which discusses how the various characteristics of a source can affect a recipient's acceptance of a message, in the context of human communication. Rieh, Hilligoss and other explore important dimensions of credibility in the context of social interactions [13], [17], [18], such as trustworthiness, expertise and information validity. McKnight and Kacmar [13] study a unifying framework of credibility assessment in which three distinct levels of credibility are discussed: construct, heuristics, and interaction. Their work is in the context of assessing the credibility of websites as sources of information.

Wright and Laskey [19] discuss how to tackle fusion of credible information. They present a weighting based, probabilistic model to compute uncertain information credibility from diverse sources. Several techniques are combined with this model, like prior information, evidence when available and opportunities for learning from data.

Sometimes, the terms credibility and trust are used synonymously. However, they are distinct notions: while trust refers to beliefs and behaviors associated with the acceptance of risk, credibility refers to the believability of a source, and a believable source may or may not result in associated trusting behaviors [17].

Finally, there is a body of work that has examined processes and propagation of credible information. Corroboration as a process of credibility assessment is discussed in [20]. Proximity, both geographic and social, and its role in credibility assessment is discussed in [5]: our role of geographic distance as a measure of credibility is related to this discussion. Saavedra et al. [21] explore the dynamics and the emergence of synchronicity in decision-making when traders use corroboration as a mechanism for trading decisions.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have explored the design space of algorithms for a new problem, optimizing pull corroboration in an emerging application area, social swarming. We have proposed optimal special-case algorithms, computationally efficient approximations, and decentralized optimal stochastic variants. However, our work is merely an initial foray into a broad and unexplored space, with several directions for future work: increasing credibility and cost model realism, incorporating malice, allowing peers to relay reports, and exploring other realistic, yet efficient and near-optimal special-case solutions.

REFERENCES

- [1] M. Min, S. Reddy, K. Shilton, N. Yau, and et al. PEIR, the personal environmental impact report, *Proc. 7th ACM Mobisys*, June 2009.
- [2] E. Miluzzo, N. Lane, K. Fodor, R. Peterson, and et al. Sensing meets mobile social networks: the design, implementation and evaluation of the CenceMe application, *Proc. 6th ACM SenSys*, November 2008.
- [3] S. Kang, J. Lee, H. Jang, H. Lee, and et al. SeeMon: scalable and energy-efficient context monitoring framework for sensor-rich mobile environments. *Proc. 6th ACM MobiSys*, June 2008.
- [4] R. K. Balan, D. Gergle, M. Satyanarayanan, J. Herbsleb. Simplifying cyber foraging for mobile devices, *Proc. 5th ACM MobiSys*, June 2007.
- [5] M. T. Rivera, S. B. Soderstrom, and B. Uzzi. Dyads: Dynamics of Dyads in Social Networks: Assortative, Relational, and Proximity Mechanisms. *Annual Review of Sociology*, Vol. 36, pp 91-115, Jun 2010.
- [6] M. Ra, J. Paek, A. B. Sharma, R. Govindan, M. H. Krieger, and M. J. Neely. Energy-Delay Tradeoffs in Smartphone Applications. *Proc. 8th ACM Mobisys*, June 2010.
- [7] S. Martello and P. Toth. *Knapsack Problems: Algorithms and Computer Implementations*. Wiley: Chichester, England, 1990.
- [8] R.K. Ahuja, T.L. Magnanti, J.B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [9] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking*, vol. 1, no. 1, pp. 1-149, 2006.
- [10] M. J. Neely. Energy optimal control for time varying wireless networks. *IEEE Transactions on Information Theory*, vol. 52, no. 7, pp. 2915-2934, July 2006.
- [11] M. J. Neely. *Dynamic Power Allocation and Routing for Satellite and Wireless Networks with Time Varying Channels*. PhD thesis, Massachusetts Institute of Technology, LIDS, 2003.
- [12] M. J. Neely. Stability and capacity regions for discrete time queueing networks. *ArXiv Technical Report: arXiv:1003.3396v1*, March 2010.
- [13] B. Hilligoss and S. Rieh. Developing a Unifying Framework of Credibility Assessment: Construct, heuristics, and Interaction in Context. *Information Processing and Management*, Vol. 44, 1467-1484, 2008.
- [14] C. I. Hovland and W. Weiss. The Influence of Source Credibility on Communication Effectiveness. *Public Opinion Quarterly*, Vol. 15, pp 635-650, 1951.
- [15] D. H. McKnight and C. J. Kacmar. Factors and Effects of Information Credibility. *Proc. ICEC'07*, Minneapolis, 2007.
- [16] B. J. Fogg and H. Tseng. The Elements of Computer Credibility. *Proc. CHI'99*, 2004.
- [17] S. Rieh and D. Danielson. Credibility Models for Multi-Source Fusion. *Annual Review of Information Science and Technology*, Vol. 41, pp 307-364, 2007.
- [18] P. Chen. Information Credibility Assessment and Meta Data Modeling in Integrating Heterogeneous Data Sources. *Air Force Research Laboratory Technical Report AFRL-IF-RS-TR-2002-298*, 2002.
- [19] Wright, E. and Laskey, K. Trust in Virtual Teams: Towards an Integrative Model of Trust Formation. *9th International Conference on Information Fusion*, Florence, Italy, 2006.
- [20] R. B. Cialdini. The Science of Persuasion. *Scientific American Mind*, Jan 2004.
- [21] S. Saavedra, B. Uzzi, and K. Hagerty. Synchronicity and the collective genius of profitable day traders. Working Paper, Kellogg School of Management, 2010.
- [22] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.

- [23] Y. Hung, A. Dennis and L. Robert. Trust in Virtual Teams: Towards an Integrative Model of Trust Formation. *37th International Conference on System Sciences*, 2004.
- [24] Grabner-Krauter, S., Kaluscha, E., Fladnitzer, M. Perspectives of Online Trust and Similar Constructs: A Conceptual Clarification. *8th International Conference on Electronic Commerce (ICEC)*, New Brunswick, Canada, 2006.
- [25] Cho, J. and Swami, A. Towards Trust-Based Cognitive Networks: A Survey of Trust Management for Mobile Ad Hoc Networks. *14th International Command and Control Research and Technology Symposium*, Washington, DC, 2009.