

Efficient Coordinated Power Distribution on Private Infrastructure

Andrew Perrault
Department of Computer Science
University of Toronto, Toronto, CANADA
perrault@cs.toronto.edu

Craig Boutilier
Department of Computer Science
University of Toronto, Toronto, CANADA
cebly@cs.toronto.edu

ABSTRACT

Current power distribution network design makes it attractive for agents to generate their own power (distributed generation) and to construct private infrastructure (e.g., distribution lines) to exchange power without using the main public grid. We show that such private transactions may increase overall network load because of increased transmission distances, thus increasing resistive losses. We present a coordination scheme for the centralized control of private infrastructure that satisfies participation constraints and budget balance. Experiments show that our scheme reduces distribution losses by 4-5% when there are only a constant number of private lines and by 55%-60% when the number of private lines is proportional to the number of agents.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence

General Terms

Agents, Multi-Agent Systems, Economics

Keywords

Energy and Emissions, Incentives for Cooperation, Game Theory

1. INTRODUCTION

The *smart grid* is roughly defined as an augmented electrical grid that gathers information about its own operation to automatically improve efficiency and reliability. Utility companies have been the primary driver of development of the smart grid thus far, seeing an opportunity to replace monitoring tasks undertaken by people with centralized electronic sensing and control. Full realization of the smart grid will require intelligent control and incentive schemes to make the best use of information so gathered [15].

In recent years, it has become increasingly easy for “household” consumers to generate electricity at cheaper cost and with lower emissions using renewable sources—most notably wind and solar power. Installation capacity is often higher in rural areas where space is available. Groups of

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

such consumers can often limit their dependence on the public grid by exploiting local generation, e.g., in the form of *microgrids* [10]. Microgrids are groups of generators and consumers connected to the larger grid at a single point, which allows decoupling in the event of grid failures. For instance, consider a rural area in a developing country where a significant amount of power is produced by wind farms, solar panels or biofuels. These tend to be located nearer to farms than to urban centers, so local generators may strike deals with nearby consumers (e.g., farms) to construct private lines and exchange power outside the public grid. Indeed, due to fixed rate schedules that mediate transactions with the public grid, it is often more profitable for them to sell directly to their neighbors than to the grid.

Unfortunately, as we will see below, such transactions can reduce overall efficiency. The dynamic above creates tension between the publicly-managed grid and the private agents it serves. Though the public grid may wish to construct additional distribution capacity, especially in the presence of increasingly distributed, unreliable generation, we show that the benefit of doing so is limited by the degree of private control of distribution lines allowed. The network in Fig. 1 (explained further in Sec. 3) illustrates this. Integrating microgrids and designing control and incentive schemes that mitigate this potential inefficiency are important challenges in smart grid design.

In this work, we develop a routing scheme that allows agents in a local microgrid to coordinate private distribution with the publicly-owned utility. Our scheme achieves optimal distribution at minimum cost while satisfying participation constraints—each agent is at least as well off as in the uncoordinated regime. Our scheme can also be used as a basis for determining cost-sharing of new local generation and distribution infrastructure among agents by quantifying their individual benefits assuming optimal usage.

This rest of this paper is organized as follows. In Sec. 2, we briefly describe our setting, outline our model of the problem, and discuss related work. In Sec. 3, we describe several models of agent behavior and the corresponding incentives. In Sec. 4, we present an incentive/routing scheme that coordinates agents on the private network with the public network. Experiments in Sec. 5 compare the performance of our system under various models of agent behavior. We present conclusions and future directions in Sec. 6.

2. SETTING AND NETWORK MODEL

Setting: Current electricity pricing regimes in many jurisdictions require that electric utilities report the expected demand of their consumers, and generators report their ex-




Figure 1: A network with private and public links. Circles represent net generators and squares net consumers. Node T represents the connection to the outside power grid. Solid lines are publicly-owned links and dashed lines private links. We assume that public and private links have roughly the same voltage and resistance. Q_i denotes the supply of agent A_i .

pected production, to an *independent system operator (ISO)* [7]. ISO-controlled markets are an example of a *mechanism* where a central protocol maps declarations by market participants to transactions and payments. The ISO typically solves a linear relaxation of the *distribution problem*, telling generators how much power to produce in a way that neglects the non-linear physics of the system. This relaxation is more accurate in relatively uncongested networks or those with short distribution lines, and can be solved efficiently and transparently. But a tension exists between easy-to-calculate, transparent prices and prices that incentivize optimal behavior. The ISO solution is relatively complex, but its implementation can be analyzed and simulated by market participants to determine how to most advantageously represent their supply, demand or production costs.

The generation and distribution context we analyze is speculative in certain aspects, as it describes a system that is still under debate by policymakers. At a high level, we assume that small producers and consumers will have significant control over trading decisions in future systems, since increasing the efficiency of electricity markets will likely be accomplished through the deregulation made possible by smart grid technology and fine-grained control of routing decisions. We also note that nearly all research in power distribution uses some approximate model of electrical physics since exact power flow is highly non-convex. We simplify the physics to focus on the dynamics of electricity exchange between pairs of agents. Our most significant simplification is neglecting sympathetic flows (see below), but we believe they can be incorporated into our model as well.

Network Model: We first introduce the notation needed to describe our model. The network $G = (V, E)$ is composed of two parts, a *public network* $G_{pub} = (V, E_{pub})$ that is connected and a *private network* $G_{priv} = (V_{priv}, E_{priv})$ where $V_{priv} \subseteq V$. A single node T represents the link to the outside grid, which will buy or sell arbitrary amounts of electricity; it is not connected to any agent via the private network. Let $f_{u,v}$ denote the power sent from agent u to agent v over edge (u, v) , and let $f'_{u,v}$ denote the power received by agent v over edge (u, v) after losses are deducted. These variables will be constrained to be non-negative. The $f_{u,v}$ variables are used to formulate the problem as an optimization, but the solutions will always have at least one of $f_{u,v}$ and $f_{v,u}$ equal to zero, since power may only flow in one direction at a given time.

Each agent A_v (except T) has a *utility function* $U_v(x_v)$ for any allocation x_v of power representing their net consumption. Because demand for electricity is usually highly inelastic, we represent a utility function by a single constant Q_v ,

the agent’s, steady-state (long-term) supply or demand for power. We will assume that the variable cost of distributed generation sources is near zero (realistically, since they are renewables).¹ Positive Q_v reflects a surplus of electricity, in which case we assume $U_v(x_v)$ is very steep for $x_v < 0$ and equal to zero thereafter. Negative Q_v reflects net demand, where $U_v(x_v)$ is very steep for $x_v < -Q_v$ and zero after. We will only consider settings in which all demands can be feasibly satisfied, so we treat these utility functions as if they were *constraints* on the final solution. These assumptions about utility functions simplify the analysis significantly.

Fig. 1 illustrates the model: each circle is a net generator and each square a net consumer, with Q_v being net power generated by agent A_v . T represents the connection to the public grid, while public and private lines are solid and dashed lines, respectively.

We consider several models of agent behavior below corresponding to different degrees of coordination and incentive structures. Each model, together with an agent’s connections, will determine that agent’s strategy space. All models below share a common simplified physics model:

- Flow of electricity must be conserved: flow into any v must equal flow out of v less the agent’s supply Q_v .
- *Resistive losses* occur when power is transmitted between two nodes. Energy lost to heat and radiation is proportional to the *square* of the amount transmitted, and depends on the voltage and resistance of a particular link. Let $R_{u,v}$ be the resistance and $U_{u,v}$ the voltage on edge (u, v) . If u sends $f_{u,v}$ units of power over the line (u, v) , the amount v receives from u is:

$$f'_{u,v} = f_{u,v} - \frac{f_{u,v}^2 R_{u,v}}{U_{u,v}^2} \quad (1)$$

Note that $R_{u,v}$ is proportional to the length of (u, v) . Resistive losses are assumed to be negligible in many simplified power flow models, but their inclusion—even though they are relatively small (about 7% avg.)—changes the dynamics of the distribution game.

- Transformer losses occur when electricity is “stepped up/down” between high public grid voltages and low local grid voltages at T . This loss is a constant fraction β of the power converted.
- Electricity can be disposed of at any network node.²

From a physics standpoint, our simplified model treats alternating current (AC), the choice of most power systems worldwide, as direct current (DC), which is less complicated to deal with, and is a common simplification in power systems research (flow prediction in AC networks is often non-convex) [11]. We also neglect Kirchhoff’s voltage law, that causes power to flow sympathetically on loops. Below we introduce an energy-minimization problem that will approximate physics-induced flow on our simplified model.

¹There are obvious limits to the assumption of total demand inelasticity in electricity consumption, with studies restricted to small price ranges. The problem of making electricity demand more sensitive to cost, whether through price or other means (e.g., visual feedback) is an important one that is beyond the scope of our work.

²This assumption simplifies computation. In real systems, free disposal may be used to compute payments and more accurate computation to predict flows.

To justify our physical model, we observe that it gives an upper bound on the efficiency of true physics-based routing; in particular, AC power introduces interference when currents are out of phase, and Kirchhoff’s law introduces “phantom” power flows that are not useful for achieving the routing objective. We also note that any physics model can be substituted into the analysis without affecting calculations in other areas. We use an “optimistic” high-level physics model to show that the “goals of physics” are not the same as those of a grid operator; in particular, physics-based routing does not take into account the cost of generation from a particular source. High-cost power can be overused if it is close to its consumers, even though using far-away low-cost power, despite incurring higher losses, may be more cost effective and have lower emissions.

To see how private routing can affect the network, consider Fig. 1. Since the grid’s buy and sell prices are different, A_5 and A_6 will exchange power over the privately owned dashed edge, at a mutually beneficial price between the grid buy/sell prices. The public grid must then send power over a long distance from T to A_7 . If private edge (5, 6) did not exist, power would only need to be sent from T to A_5 , which is closer. Due to resistive losses, the existence of the private edge between A_5 and A_6 combined with rational behavior of A_5 and A_6 *decreases* the social welfare of the resulting routing equilibrium.

Below we outline different ways in which *coalitions* of agents can coordinate their behavior to maximize their collective utility, assuming they can divide any surplus generated appropriately. The public network is centrally managed, but routing on any private link can be arbitrarily controlled. We assume agents in separate coalitions do not trade with each other on the private network (otherwise, they will be taken as part of the same coalition).

Two coordination issues arise in this setting. The first is finding optimal coordinated strategies, or a *routing scheme*, for the *grand coalition* of all agents: complete coordination will ensure all agent demands are met at minimum cost. The second is finding an *incentive scheme*, or payments, that stabilize the grand coalition by aligning their interests to induce minimum cost flow on the network. We first address the former, assuming that agents can find stabilizing payments.

Related Work: Our model can be viewed as a cooperative game [14], though because of potential externalities imposed by one coalition on another—if two coalitions use the same line, losses are greater than if only one did—standard characteristic function representations do not apply directly. If not for this, our model could be viewed as a *market game* [20], where supply/demand correspond to initial endowments and trading is restricted to occur between agents in the same coalition. The core of a market game is non-empty if agent utilities are concave and no losses occur in trading. Extensions that accommodate loss due to trading (e.g., transaction costs) [17] are needed in our model since different links have different losses that vary with load. Models of prices sensitive to physical constraints have been proposed [5], but place a large “choice” burden on market participants and ignore transactions outside the market.

Previous work in communication networks has produced edge-based market-clearing prices that induce Nash equilibria in routing games, for control in networks without delays [9, 8] and monopolist/oligopolist-control in restricted classes

of networks with convex delays [1, 2]. The quality of such networks has also been studied [4, 6]. These latter models reflect the same type of congestion-induced externalities that arise in our model, but do not offer the generality of analysis we need. Our model of the physics of distribution was inspired primarily by [18, 24].

Smart grid research has largely focused on hardware development, but has been an area of increasing concentration in AI [15], primarily addressing coordination among market-facing agents. One goal of this research is to shape demand to reduce maximum loads and react to fluctuations in supply availability (e.g., by shifting some peak demand to off-peak hours). Peak demand is a critical driver of supply costs (typically provided by finely adjustable, but expensive oil or gas turbines). Large-scale distributed battery capacity can address this problem (e.g., using the batteries of idle electric cars) but specific coordination algorithms have a significant effect on efficiency [23]. Another approach is to coordinate battery charging and use [3]. Our approach differs from the perspective laid out by Ramchurn et al. [15] as physical network constraints make distributed computation difficult. We differ from Alam et al. [3] by focusing on efficiency in systems where the infrastructure is owned by different agents and losses occur in distribution.

Recent work has looked at coalitions of consumers with complementary demands, which are cheaper to supply because of their flatter load profiles. Incentives for coalition formation have been investigated [21]. To relieve consumers from constant usage decisions based on supply availability, intelligent agents can manage usage and storage based on user preferences [22]. Similar coordination issues exist on the supply-side of the market; for instance, the unpredictability of wind power means that *coalitions* of generators can more easily negotiate contracts with ISOs [16].

3. MODELS OF AGENT BEHAVIOR

We compare optimal agent routing behavior under four different models of cooperation. In addition to a mathematical specification for each model, we provide a description of the flows it induces on the network in Fig. 1. Induced flows in each model are computed sequentially in two stages. In the *private* stage, agents on the private network agree to exchange power. In the *public* stage, routing on the public network occurs given net demands after the private stage, and any payments are made between the grid and the agents.

The Ad Hoc Model. The *ad hoc model* simulates system behavior when agents make decisions in a local, myopic manner without coordination or cooperation. Agents will only ever trade power with their neighbors, and they do so in order of closeness. Each trade is made to maximize the joint utility of the two traders. This model probably best reflects *status quo* behavior because private lines are unregulated, and there are few of them. It is unlikely that agents that are very distant from each other will be able to discover that it would be beneficial to build a private line between them. Although the random network generation model we use in experiments below does not take distance into account, this behavior model favors the use of lines between agents that are closer to each other. The routing on the public network is induced by physics, which aligns with the primarily physics-based routing that occurs on the public grid. To compute exchanges on the private network, we order private edges by increasing length—agents that are close together are more likely to transact because they are generally more familiar with each other. For each $(u, v) \in E_{priv}$,

the non-negative flow on that edge is set such that the total utility of agents u and v is maximized:

$$\max_{f_{u,v}} \begin{cases} p_s(Q_u - f_{u,v}) & \text{if } Q_u - f_{u,v} \geq 0 \\ -p_b(Q_u - f_{u,v}) & \text{if } Q_u - f_{u,v} < 0 \end{cases} + \begin{cases} p_s(Q_v + f'_{u,v}) & \text{if } Q_v + f'_{u,v} \geq 0 \\ -p_b(Q_v + f'_{u,v}) & \text{if } Q_v + f'_{u,v} < 0 \end{cases} \quad (2)$$

where p_b and p_s represent the grid buy and sell prices. We then update the demands of these two nodes by setting $Q_u^{new} = Q_u^{old} - f_{u,v}$ and $Q_v^{new} = Q_v^{old} + f'_{u,v}$. We continue through the ordered sequence of edges until no profitable trades remain. Note that two agents only exchange if one is a net generator and the other a net consumer, and agents only trade with their immediate neighbors in G_{priv} .

After these local transactions have occurred, remaining demand is met using physics-based flow. This occurs without the agents choosing a strategy because they are unable to affect the flow on infrastructure that they do not own. Physics-based routing uses the physics model described in the previous section. We minimize distribution losses by solving the following optimization:

$$\min_{f_{u,v}, \forall (u,v) \in E_{pub}} \sum_{(u,v) \in E_{pub}} \frac{f_{u,v}^2 R_{u,v}}{U_{u,v}^2} + \frac{\beta}{1-\beta} \sum_{v:(T,v) \in E_{pub}} f_{T,v} + \beta \sum_{v:(v,T) \in E_{pub}} f'_{v,T} \quad (3)$$

subject to, for each node $v \in V$:

$$Q'_v + \sum_{u:(u,v) \in E} f'_{u,v} \geq \sum_{u:(v,u) \in E} f_{v,u} \quad (4)$$

where Q'_v (in this case, Q_v^{new}) is the supply of node v induced by the flows from private trading. The first term of Objective 3 represents the resistive losses incurred on every edge in the public network, while the second and third terms represent losses due to voltage step down/up at node T . Constraint 4 requires that flow be conserved at every node in the network.

We observe the following induced flows on the network in Fig. 1 under this ad hoc model. On the upper branch, A_1 sends power to A_2 , but the private edge (2,3) remains unused because A_1 is unable to trade with A_3 in the private stage. This causes public edge (2,3) to carry more power than necessary, incurring higher losses. On the middle branch, the cost of supplying A_4 is higher than necessary because physics-induced routing equalizes the losses on $(T,4)$ and $(1,4)$ —wasting power from A_1 —as long as losses on $(1,4)$ exceed β . On the lower branch, A_6 sends its power to A_5 , leaving A_7 to receive power from T . This is inefficient because A_5 is closer to T than is A_7 .

The Private Self-Interest Model. The *private self-interest model* reflects agents on the private network cooperating to maximize their aggregate utility under fixed public network pricing. The flow on the public network is determined by physics, analogous to that in the ad hoc model (minimizing distribution losses under basic physical constraints). Agents on the private network solve the following optimization:

$$\min_{f_{u,v}: (u,v) \in E_{priv}} \sum_{v \in V_{priv}} \begin{cases} -p_s(Q_v + x'_v) & \text{if } Q_v + x'_v \geq 0 \\ p_b(Q_v + x'_v) & \text{if } Q_v + x'_v < 0 \end{cases} \quad (5)$$

where

$$x'_v = \sum_{u:(u,v) \in E_{priv}} f'_{u,v} - \sum_{u:(v,u) \in E_{priv}} f_{v,u} \quad (6)$$

Objective 5 represents the total cost to nodes on the private network to satisfy remaining demand from (or sell surplus to) the public grid after all (optimal) private transactions are made. After the private agents have minimized their liability to the grid, physics governs flow on the public network as in the ad hoc model (Objective 3 subject to Constraint 4).

In this model, the flow on the middle and lower branches of Fig. 1 is the same as in the ad hoc model. On the upper branch, A_1 sends power to A_3 during the private stage (exchanges between non-adjacent nodes is now permitted). However, this does not reduce the cost of power drawn from the public network; it merely shifts load from public to private edges, leaving public edges virtually unused if there is little production on the private network.

By solving the optimization, agents on the private network maximize their aggregate utility. However, they still must divide up the surplus in a way so that none of their agents defect. We discuss this problem in Sec. 4.

The Cooperative Model. In the *cooperative model*, agents on the private network distribute electricity so that the load on the public network is minimized (again, flow on the public network is determined by physics). The agents solve a joint optimization that minimizes the overall cost to supply electricity to the entire network by minimizing the amount of externally generated power consumed by the *entire* set of agents (according to our assumption that all locally generated power is from renewable sources and thus has near zero marginal cost). It requires suitable incentives to induce cooperation, as we discuss in Sec. 4. Minimizing the overall draw from the grid requires arranging flows on the private network that solve the following optimization:

$$\min_{f_{u,v}: (u,v) \in E_{priv}} \frac{1}{1-\beta} \sum_{v:(T,v) \in E} f_{T,v} - (1-\beta) \sum_{v:(v,T) \in E} f'_{v,T} \quad (7)$$

subject to the usual flow and physical constraints, Eqs. 4 and 3. The terms in Objective 7 represent the total power flow into and out of the network through T , after accounting for step-up/down losses at the transformer. This optimization is more difficult than the others since flows on the public network are set using a separate objective (physics-based energy minimization). This computation is explained below.

The flows induced by the cooperative model in Fig. 1 are quite different from those in the previous models. Flows on the upper and lower branches resolve the earlier inefficiencies: the upper flow from A_1 to A_2 and A_3 uses the public and private links equally to minimize resistive losses, and on the lower branch, power from A_6 flows to A_7 instead of A_5 . Despite the improvement, flow on the middle branch remains inefficient due to physics-based routing, causing A_4 to draw power from both T and A_1 .

The Integrated Model. In the *integrated model*, we allow the flow on all edges, both private and public, to be controlled by the local agents, relaxing the constraint that flow on public edges be governed by the simple physics model (which minimizes losses). This can be viewed as “ideal” behavior for agents on a network that have arbitrarily precise routing equipment and some degree of control over (local) distribution on public lines. This reflects the most effective use of resources if lines are expensive. It requires solution of

the following optimization:

$$\min_{f_{u,v}:(u,v) \in E} \frac{1}{1-\beta} \sum_{v:(T,v) \in E} f_{T,v} - (1-\beta) \sum_{v:(v,T) \in E} f'_{v,T} \quad (8)$$

subject to constraints Eq. 4. The objective is identical to that in the cooperative case (Eq. 7) except that *all flows* are controllable, not just those on private links.

The flows induced by the integrated model in Fig. 1 are the same as those of the cooperative model on the upper and lower branches. However, because no physics-based flows occur, the inefficiency on the middle branch disappears, and A_4 will draw power only from A_1 .

Computation. Except for the Cooperative model, the optimization problems above are all convex, quadratically-constrained quadratic programs (QCQP) that can be solved by CPLEX or other off-the-shelf optimizers. In the cooperative case, we approximate the solution by first computing flows on the private network as in the integrated model, and then (approximately) minimizing the objective by computing the physics-induced flows on the public network. The result is a feasible solution whose quality, in our experiments, is within a few percent of that of Integrated (which in turn offers a lower bound on Cooperative).

4. A SIMPLE INCENTIVE SCHEME

While the methods above compute optimal coordinated behavior, we must also ensure incentives can be put in place such that all agents effectively cooperate by actually adopting the strategy prescribed by the optimal solution. The aim of these prices is to induce the cooperative or integrated behavior described in the previous section. Which of the two models can be used in practice depends on what kind of routing equipment is available on the network.

Equilibrium Prices. We first present a pricing scheme that aligns the interests of each agent on the private network with socially optimal behavior. This induces self-interested agents to adopt the cooperative solution (or the integrated solution if public links are controlled locally). Our technique is relatively standard in network models [19, 13, 9, 1]: we compute agent-specific buy/sell prices using dual prices obtained from the optimization. These induce individual agents to implement the globally optimal solution by maximizing their own net utility at these prices.

We first solve the integrated optimization Eq. 8. By the first-order optimality conditions, there is a set of dual variables $\mu_v \geq 0$ for each agent $v \in V$ such that: (a) for each edge (u, v) , where $u, v \neq T$:

$$-\mu_v \left(1 - \frac{2f_{u,v}R_{u,v}}{U_{u,v}^2} \right) + \mu_u = 0; \quad (9)$$

(b) for each edge (T, v) :

$$\frac{1}{1-\beta} - \mu_v \left(1 - \frac{2f_{T,v}R_{T,v}}{U_{T,v}^2} \right) = 0; \quad (10)$$

and (c) for each edge (v, T) :

$$-(1-\beta) \left(1 - \frac{2f_{v,T}R_{v,T}}{U_{v,T}^2} \right) + \mu_v = 0. \quad (11)$$

The values of these dual variables “price” the direct transmission to and from each agent. We charge each agent v :

- μ_v for each unit sent to $u \neq T$ and $-\mu_u$ for each unit received (after losses) by $u \neq T$.
- $-\mu_u$ for each unit sent from $u \neq T$ and μ_v for each unit received (after losses) from $u \neq T$.

Assume these prices are fixed. Each v in the local network has the following “personal” optimization to maximize its utility for the (net) power received minus the price paid:

$$\begin{aligned} & \max_{f_{u,v}:(u,v) \in E} U_v \left(Q_v + \sum_{u:(u,v) \in E} f'_{u,v} - \sum_{u:(v,u) \in E} f_{v,u} \right) \\ & + \sum_{u:(v,u) \in E} (\mu_v f_{v,u} - \mu_u f'_{v,u}) + \sum_{u:(u,v) \in E} (-\mu_u f_{u,v} + \mu_v f'_{u,v}) \end{aligned}$$

The first-order optimality conditions for this problem are a subset of those for global optimization. Thus, the optimal flow for any agent at the given prices coincides precisely with the globally optimal flows. Furthermore, no *coalition of agents* can improve their aggregate utility by changing their strategies at these prices. In other words, these prices induce a *strong Nash equilibrium* that coincides with the globally optimal outcome. However, these prices have some weaknesses. We have no guarantee about the total amount paid to the mechanism (ideally, it is zero). We also cannot know if the agents would participate if we do not know their previous contracts. We address these issues further below.

Core Stability. If the prices that arise as the by-product of optimization are taken as given, they induce strongly stable optimal behavior in utility-maximizing agents. However, if the market coordinator has no binding power to force agents to accept these prices, some may refuse to “enter this market” if they had prior local arrangements that gave them higher utility. For example, in a network with a single private edge between a generator and a consumer, these two agents will likely benefit greatly by using this edge, and be unwilling to agree to a change in pricing without inducement. Furthermore, while coalitional defections cannot be beneficial *given these fixed prices*, the allocation induced by these prices cannot generally be supported if one or more agents refuse to participate. Thus, we need to consider payment schemes that prevent such defection.

Superficially our models seem to fit within the framework of market games, albeit with variable “transaction” costs induced by distance and load. Unfortunately, there is no natural characteristic function representation. The value that one coalition of cooperating agents receives from a given strategy can be altered by the actions of another: if a second coalition routes power over the same lines, the higher load increases power loss quadratically. Research on market games [20] shows that core transfers exist in a variety of exchanges by using dual variables of allocation optimization as prices, as we do above, making core transfers easy to compute. However, our model is not reducible to market games due to interaction effects. These require that, when reasoning about potential defections from the grand coalition, we should consider Nash equilibria that emerge when treating different coalitions as players.

We have not yet been able to prove that transfers exist that support a strong Nash equilibrium (or loosely, a *core imputation*). We strongly suspect that this is so: in *all instances* tested empirically, we are able to compute stability-inducing transfers. This remains an important topic of re-

search. We note that if the market coordinator (system operator) knows of existing arrangements between agents prior to joining the coordinated local network, participation constraints that induce them to participate can be easily addressed. If the agents were to agree to participate in the pricing scheme, the social welfare of the system will increase by a non-negative amount, since the current distribution schedule is a feasible option in the global optimization. The operator can distribute this surplus to agents on the private network, guaranteeing that each receives utility at least equal to that derived prior to coordination.

5. EXPERIMENTS

We ran experiments to test the impact of using each of the four models of agent behavior in a variety of scenarios. We use both random and existing models of public networks; but since private infrastructure is rare (except for self-contained microgrids with single-point connections to the public grid), we make some general assumptions about private-network structure and test various parameterizations. Some reflect current practice, while others reflect future scenarios with increased usage of local, private generation and infrastructure. Our proxy for social welfare is *distribution loss*, i.e., the amount of power originating in the public grid not consumed by agents in the local network (assuming all agent demands are satisfied). This correlates linearly with social welfare assuming local producers have linear cost curves for generation.³ All local networks in our experiments (realistically) consume more power than they produce and have average transmission and distribution loss of around 7% (with std. dev. of 2.5%) when no private infrastructure or distributed generation is assumed: this is the current U.S. average (with most losses due to distribution).

Our first experiments use random public networks: 100 agents are distributed uniformly over a square grid of size either 100×100 km (low density) or 1×1 km (high density). We generate a distance-minimizing spanning tree connecting these agents (and T), reflecting how most public distribution networks are constructed. Links between agents have voltage of 22kV, and links to T are 50 kV (higher load links are typically higher voltage). All edges have 0.2 ohms of resistance per km. Transformer losses at T are given by $\beta = 0.02$ (this determines the theoretical lower limit on distribution loss for a net-consumer network, as all publicly-supplied power must pass through T). We use a 10% gap between the grid buy price p_b and sell price p_s .

Agent demands are distributed normally $N(\mu, \sigma)$, with σ varying with the amount of private generation: we set $\sigma = \mu/2$ (low private generation), $3\mu/4$ (medium), and $5\mu/4$ (high), meaning that we expect 2.5%, 9%, and 21% of agents (resp.) to be net producers. The private network is generated by sampling an Erdős-Renyi random graph, with p determining edge density: each edge is included with probability p . We use edge densities of 0 (i.e., no private network), $\frac{1}{n^2}$, $\frac{1}{n}$, and $\frac{1}{\log(n)}$ (the last representing heavy investment in private infrastructure).

Fig. 2 illustrates our experimental setup, 40 agents for clarity: this network has 22 private and 39 public edges and high distributed generation. Power flows mainly on the public network in the Ad Hoc and Private Self-interest models,

³Since cost curves are typically convex, our results *underestimate* the true improvement in social welfare.

but the private network is used more in the Cooperative and Integrated models. In this scenario, distribution losses are 11.3% in Ad Hoc, 11.4% in Private Self-interest and 3.8% in Cooperative and Integrated.

Tables 1 and 2 summarize our results on random public networks, with each data point showing average distribution loss over 1000 random instances.⁴ At a high level, we see that Cooperative and Integrated provide significantly lower losses relative to Ad hoc and Private Self-interest in nearly all cases. We see relatively small amounts of savings at private edge densities of 0 and $\frac{1}{n^2}$ and relatively large amounts at $\frac{1}{n}$ and $\frac{1}{\log(n)}$. Increased private network density correlates very strongly with savings in the more coordinated models: losses are roughly 60% less than in Ad hoc and Private Self-interest when edge density is $\frac{1}{n}$ or greater (in both the low and high node density models). With edge density $\frac{1}{\log(n)}$, distribution losses approach the lower bound of 2% associated with purchase of all power from the grid with no resistive (only transformer) loss.

Our second set of experiments uses the *IEEE 300-bus test system* as the public network, with the same random private networks as described above. The test system has 300 nodes and 409 edges and has balanced supply and demand due to the presence of a few large power producers. Since the system lacks spatial coordinates, we use a “spring-based” energy minimization algorithm to estimate the relative positions of the nodes. We then scale these positions in order to test both dense and sparse configurations. Table 3 shows results of 100 trials.⁵ The results are broadly similar to those above except in Private Self-interest, where losses spike to 20% at edge density $\frac{1}{n}$ and then fall to around 2.5% at density $\frac{1}{\log n}$. We discuss this difference below in the context of model performance.

Theoretically, Integrated should be more efficient than Cooperative, which in turn should outperform Ad hoc and Private Self-interest; indeed our results reflect this.⁶ No analysis yet suggests how Ad hoc and Private Self-interest models compare. In practice, they perform similarly on random public networks, but very differently in the IEEE test network. This is due to supply distribution in the test network, with just a few large producers. When a (graph) component forms at $\frac{1}{n}$ edge density containing a large producer, all demand in that component will be met; but at the expense of efficiency. At density $\frac{1}{\log(n)}$, demand is met efficiently using only the large private network. This does not happen in the random models (where private networks are net consumers).

Our results assume a fixed buy-sell price gap of 10%. The smaller this difference, the less trading occurs in the Ad hoc and Private Self-interest models. A small gap mediates the effect of bad routing decisions on the private network, but also causes less usage of private edges, increasing load on public infrastructure. Tables 1 and 2 show a mild increase in losses as more distributed generation is deployed. This is

⁴Std. dev. is 2.5-3% for losses in the 7% range, 0.5% around 3%, and 0.01% around 2%. All differences between Ad hoc/Priv. self-int. vs. Coop./Integrated with non-zero edge density are significant at $p < 0.05$.

⁵Std. dev. is about 13% in Priv. at density $\frac{1}{n}$, 1% for Ad hoc at $\frac{1}{n}$ and $\frac{1}{\log(n)}$, and less than 0.2% otherwise.

⁶With low node density, low distributed generation and $\frac{1}{n}$ edge density, Integrated performs slightly worse than Cooperative (due to build up of rounding error).




Figure 2: An example network under different models of agent cooperation. From left, Ad Hoc, Private Self-interest, Cooperative and Integrated. Net consumer agents are white, net suppliers are blue, and the network connection is black. Edge darkness represents amount of flow. Public network edges are thicker than private network edges.

not surprising since this also introduces higher demands in our model (we keep *net local demand* constant)—load becomes more focused around large producers and consumers, causing large losses at these points. Average loss in the zero private edge case is 7.19% in low distributed generation and 7.29% in high distributed generation. With $\frac{1}{\log(n)}$ edge density, average loss is 4.69% in low distributed generation and 4.73% in high distributed generation.

Fig. 3 compares losses for Ad Hoc and Cooperative as load on the network increases in a scenario with medium distributed generation and density $\frac{1}{n}$ using random public networks. Losses increase much more rapidly in the Ad Hoc model, albeit with high variance due to the many random parameters of each trial. Our analysis focuses on low-loss scenarios, which are those primarily seen in the developed world; but it is important to note that the gap in model performance increases rapidly as load approaches the limits of the distribution system.

Computationally, networks of this size were easy to solve, with the convex quadratically-constrained quadratic programs requiring well under one second. While the microgrid setting is not problematic (most networks have less than 150 nodes), scaling issues might arise in larger networks. We tested larger settings of up to 300 nodes and 90,000 edges, and were able to solve them in under an hour on a 12x2.6 Ghz system. With knowledge of prior payments, or with the ability to impose a pricing scheme, computing stabilizing payments in this context is also easy since it only requires the values of the dual variables. Without this, stabilizing payments may require approximation.

Amt. of dist. gen.	Model	Edge density on priv. net.			
		0	$\frac{1}{n^2}$	$\frac{1}{n}$	$\frac{1}{\log(n)}$
Low	Ad hoc	7.14	7.25	7.05	7.13
	Priv. self-int.	7.14	7.25	7.05	7.13
	Coop.	7.14	6.95	2.91	2.16
	Integrated	7.14	6.95	2.92	2.16
Med.	Ad hoc	7.28	7.19	7.15	7.20
	Priv. self-int.	7.28	7.19	7.15	7.20
	Coop.	7.28	6.92	2.93	2.16
	Integrated	7.28	6.91	2.93	2.16
High	Ad hoc	7.32	7.26	7.28	7.22
	Priv. self-int.	7.32	7.26	7.27	7.20
	Coop.	7.32	6.93	2.95	2.16
	Integrated	7.26	6.87	2.94	2.16

Table 1: Avg. distribution loss as a percentage of net demand in a low (node) density network (1 agent per sq. kilometer).

6. CONCLUSIONS

We have presented a routing scheme that coordinates control of private infrastructure with the public grid, improving social welfare while satisfying participation constraints

Amt. of dist. gen.	Model	Edge density on priv. net.			
		0	$\frac{1}{n^2}$	$\frac{1}{n}$	$\frac{1}{\log(n)}$
Low	Ad hoc	7.25	7.28	7.31	7.33
	Priv. self-int.	7.25	7.28	7.31	7.33
	Coop.	7.25	6.97	2.93	2.16
	Integrated	7.24	6.97	2.93	2.16
Med.	Ad hoc	7.22	7.39	7.40	7.08
	Priv. self-int.	7.22	7.39	7.40	7.09
	Coop.	7.22	7.05	3.00	2.16
	Integrated	7.21	7.04	3.00	2.16
High	Ad hoc	7.30	7.26	7.50	7.41
	Priv. self-int.	7.30	7.26	7.50	7.39
	Coop.	7.30	6.91	3.00	2.16
	Integrated	7.26	6.85	2.99	2.16

Table 2: Avg. distribution loss as a percentage of net demand in a high (node) density network (100 agents per sq. kilometer).

Node density	Model	Edge density on priv. net.			
		0	$\frac{1}{n^2}$	$\frac{1}{n}$	$\frac{1}{\log(n)}$
1 km x 1 km	Ad hoc	7.13	7.13	7.64	5.63
	Priv. self-int.	7.13	7.13	21.73	2.39
	Coop.	7.13	7.12	4.79	2.16
	Integrated	7.11	7.10	4.78	2.16
100 km x 100 km	Ad hoc	7.20	7.20	7.68	6.12
	Priv. self-int.	7.20	7.20	22.05	2.50
	Coop.	7.20	7.17	4.87	2.29
	Integrated	7.07	7.05	4.76	2.15

Table 3: Avg. distribution loss as a percentage of net demand in networks derived from the IEEE 300-bus test system.

for the agents who control private infrastructure if we assume that coalitions of agents have the ability to calculate stabilizing payments. We presented several models of agent behavior under different incentive conditions and tested the efficiency of various power distribution networks under different models. Using private infrastructure in cooperation with public infrastructure was shown to be quite important.

Although we were able to compute “core” stabilizing payments empirically in all instances, it remains open whether they always exist. Our cooperative game-theoretic model assumes that public grid prices are fixed, regardless of the behavior of agents on the private grid. This is reasonable from a short-term perspective if prices reflect long-term costs of supplying electricity, or when there are few private agents. Fixed prices also simplify analysis (generation cost curves can be ignored). If generation costs are convex, the problem remains straightforward (there is additional incentive for agents to cooperate, though further analysis is warranted).

We conjecture that there is a class of cooperative games that generalizes market games to allow non-independent congestion costs while retaining the property that core payments always exist. This class would include at least the power distribution game with fixed prices on the public grid




Figure 3: Scaling behavior, Ad Hoc (red) Cooperative (blue), as net demand increases (2nd-degree polynomial fit is shown).

as described above, but also would include games with other transmission losses (likely restricted to convex losses).

Our experimental results show that coordination is critical. Without it, private links can be constructed without decreasing overall network load. Indeed, building a private link may be profitable for an agent in the short term; but if many agents construct links without coordinating, the return on investment for the group may be zero.

Since our random network models have relatively small amounts of distributed generation (i.e., the networks are net consumers), the main routing objective is to *deliver* power as cheaply as possible. Although self-interested use of private lines reduces the demand served by the public grid, the savings are small compared to those created by using private lines to assist in distribution. In networks where supply and demand are balanced (e.g., the IEEE test network), the role of the grid changes to distributing electricity locally. In this case, private infrastructure can have a significant negative impact (as we see in our experiments) without coordination.

Some avenues for future work include relaxing our network formation assumptions (e.g., trying different random graph models), and using our incentive schemes to determine pricing and capital cost-sharing of new (private and public) infrastructure. Studying the mechanism design problem in this setting is also important. We assume that agent demands are known (e.g., from historical data) or that agents report their utility functions truthfully. Asking each agent for a complete utility function may be expensive, infeasible, or vulnerable to manipulation.

Acknowledgements: Perrault was supported by a QEII-GSST Scholarship. We gratefully acknowledge the support of NSERC.

7. REFERENCES

- [1] D. Acemoglu and A. Ozdaglar. Competition and efficiency in congested markets. *Mathematics of Operations Research*, 32(1):1–31, 2007.
- [2] D. Acemoglu and A. Ozdaglar. Competition in parallel-serial networks. *IEEE Journal on Selected Areas in Communications*, 25(6):1180–1192, 2007.
- [3] M. Alam, S. Ramchurn, and A. Rogers. Cooperative energy exchange for the efficient use of energy and resources in remote communities. *AAMAS-10*, pp.731–738, St. Paul, MN, 2013.
- [4] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *FOCS-04*,

- pp.295–304, Washington, DC, 2004.
- [5] H.-P. Chao and S. Peck. A market mechanism for electric power transmission. *Journal of Regulatory Economics*, 10(1):25–59, 1996.
- [6] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. In *DISC-03*, pp.347–351, New York, 2003.
- [7] W. W. Hogan. Electricity market restructuring: Reforms of reforms. *Journal of Regulatory Economics*, 21(1):103–132, 2002.
- [8] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.
- [9] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, 1998.
- [10] R. H. Lasseter, P. Piagi. Microgrid: a conceptual solution. *IEEE PES-04*, pp.4284–4290, Aachen.
- [11] J. Lavaei and S. H. Low. Zero duality gap in optimal power flow problem. *IEEE Transactions on Power Systems*, 27(1):92–107, 2012.
- [12] N. Nisan. Introduction to mechanism design (for computer scientists). In *Algorithmic Game Theory*, pp.209–241. Cambridge, 2007.
- [13] R. O’Neill, P. M. Sotkiewicz, B. Hobbs, M. Rothkopf, and W. Stewart, Jr. Efficient market-clearing prices in markets with nonconvexities. *Eur. J. Oper. Res.*, 164(1):269–285, 2005.
- [14] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, 1994.
- [15] S. Ramchurn, P. Vytelingum, A. Rogers, N. Jennings. Putting the “smarts” into the smart grid: A grand challenge for AI. *CACM*, 55(4):86–97, 2012.
- [16] V. Robu, R. Kota, G. Chalkiadakis, A. Rogers, and N. R. Jennings. Cooperative virtual power plant formation using scoring rules. *AAAI-12*, pp.1165–1166, Toronto, 2012.
- [17] J. D. Rogawski and M. Shubik. A strategic market game with transactions costs. *Mathematical Social Sciences*, 11(2):139–160, 1986.
- [18] W. Saad, Z. Han, and H. V. Poor. Coalitional game theory for cooperative micro-grid distribution networks. *IEEE ICC-11*, pp.1–5, Kyoto, 2011.
- [19] P. Samuelson. Spatial price equilibrium and linear programming. *Amer. Econ. Rev.*, 42(3):283–303, 1952.
- [20] L. S. Shapley and M. Shubik. On market games. *Journal of Economic Theory*, 1(1):9–25, 1969.
- [21] M. Vinyals, F. Bistaffa, A. Farinelli, and A. Rogers. Coalitional energy purchasing in the smart grid. *ENERGYCON-12*, pp.848–853, Florence, 2012.
- [22] P. Vytelingum, S. Ramchurn, T. Voice, A. Rogers, N. Jennings. Trading agents for the smart electricity grid. *AAMAS-10*, pp.897–904, Toronto.
- [23] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based micro-storage management for the smart grid. *AAMAS-10*, pp.39–46, Toronto, 2010.
- [24] A. J. Wood and B. F. Wollenberg. *Power Generation, Operation, and Control*. Wiley, NY, 1996.