

Negotiation of Usage Control Policies - Simply the Best?

Alexander Pretschner
Information Security
ETH Zurich, Switzerland
pretscha@inf.ethz.ch

Thomas Walter
DoCoMo Euro-Labs
Munich, Germany
walter@docomolab-euro.com

The term “negotiation” suggests that multi-step bidirectional communication takes place. In this position paper, we play the devil’s advocate and argue that (automated) policy negotiation essentially is one of the following, at least in the area of usage control. It can come down to a three-phase protocol that consists of a client request, a set of offers by the server, and the client’s choice of an offer or to abort. Policy negotiation can also consist of a client request together with acceptable conditions plus the server’s choice of one condition or to abort. In other words, negotiation of policies is a mere choice among alternatives; there is no negotiation in the intuitive sense of the word. — The goal of this position paper is to stimulate the discussion on what (automated) “policy negotiation” really is or can be.

The Merriam-Webster dictionary defines negotiation as the action or process of negotiating or being negotiated. To negotiate means *to arrange for or bring about through conference, discussion, and compromise*. Conference, discussion, and compromise are inherently interactive processes. Automated negotiation, the subject of this paper, is, among others, discussed in the context of agent technologies, trust, eCommerce, policies, and protocols. The relevant literature suggests that, in contrast to human negotiation, automated negotiation essentially boils down to a protocol with three or four phases only: (1) display of interesting products, (2) customer’s expression of his desire to buy one such product, (3) seller’s offers, and (4) customer’s choice. Alternatively, the following steps are taken: (1) display of interesting products, (2) customer’s expression of his desire to purchase plus conditions that are acceptable to the customer, (3) seller’s choice of one condition. The main problem, in the literature, seems to be *the evaluation and comparison* of buying and selling offers. Interaction, on the other hand, is reduced to a minimum. One could hence argue that the term negotiation is a misnomer in this context (we leave it to the discussion to propose a better term).

There are two possible reasons for the assumed lack of “true” interaction in automated settings. One is that we simply lack the creativity to think of *technologies* that would

implement negotiation with more than three or four steps.¹ The other reason, that we deem more compelling, is that the process of human negotiation includes elements that are unlikely to occur in automated negotiation. In particular, the element of *surprise* seems to be crucial.

The simplest form of negotiation is a seller who asks for too high a price; the customer, in turn, proposes too low a price; and the trade will be effected somewhere in the middle. This scenario seems unlikely in a machine setting, however—an automated customer would start at a price of 1, continue with 2, etc., until the seller agrees. More likely, the seller will directly state his price.

If people buy cars, negotiation takes place as follows (note there is none when buying a car over the internet). The potential customer enters a store, looks at different cars, and decides on a few that he finds interesting - brand, price, model, exterior, interior, and horse powers are to his taste. A first moment of surprise is that when confronted with several models, he realizes that huge exhaust pipes are a must, even though he never had thought about them before.

This is where negotiation on the price commences. The customer will have an approximate maximum in mind (that may be overridden for the sake of exhaust pipes), and the car dealer will have a lower bound. The dealer kicks off with, say, twice his lower bound; the customer argues that the same car is available at 40% of that price next door (regardless of whether or not this is true); the dealer mentions the lack of honor of next door’s dealer; the customer states that he dislikes pink leather seats (a surprise element) which must translate in a lower price; the dealer offers a set of winter tires; the customer asks for an extra navigation system (surprise - he has seen it in one of the cars), and the dealer does not agree. However, with a further generous discount for using cash (surprise), the car changes the owner.

Abstractly speaking, selling offers consist of a vector of elements that changes over time — new dimensions are added over time (current price including cash discount, win-

¹Note that we are not saying that the literature does not mention negotiations with multiple interactions—the contrary is the case. However, we are not aware of *implementations* of such protocols.

ter tires, navigation system), and the valuations of these dimensions change as well. Both buyer and seller have a preference function that either tells them whether or not this is a good deal, or tells them that this is a deal better or worse than what he has been offered before. The evaluation function needs to take into account the evolving dimensions; in particular, the dimensions that were never thought of before entering the dealer's store (exhaust pipes). Is this likely to happen without human intervention?

Some authors suggest to include multiple interactions in (future) negotiation protocols. It is unclear to us if there is a need for multiple interactions in the contexts of usage and access control. The schema of one party presenting a set of offers, the other party applying their evaluation function, and of eventually picking the best offer seems sufficient. This paper's goal is to initiate a discussion on what negotiation of policies is all about and that reveals where, in contrast to our arguments, negotiation with multiple interactions, seems sensible, if the situation is different in other domains and scenarios—or if negotiation is indeed a rather straightforward task. The determination of preference functions definitely is not.

Our background is that of usage control. We discuss negotiation for usage control for two reasons—it provides a domain from which we can draw examples, and it prevents us from stating arguments in a generality that makes them either trivially wrong (because one can always find a particular domain where a particular scenario seems sensible), or trivially true (because of the chosen level of abstraction).

Usage control is an extension of access control where control extends not only to who may access which data, but also to how the data may or must not be used or distributed afterward. In particular, we study usage control in the context of distributed systems that are composed of different actors, taking the roles of data providers (who gives data away) and data consumers (who request and receive data). Usage control is relevant in many areas, including privacy, DRM, management of IP and that of trade and administration secrets.

Data providers define policies that contain restrictions on the future usage of some data item. These policies are shipped together with the data, and then used to configure usage control mechanisms. Negotiation in the domain of usage control must thus at least consider the following dimensions: data, policies and enforcement mechanisms.

Data is the good that is being negotiated, for instance, an e-book, an mp3 song, or some piece of IP. *Policies* consist of access control requirements, provisional actions, and obligations. Access control requirements govern who may access the data at all. Provisional actions stipulate actions to be undertaken in-between the request for and the release of data, e.g., up-front payments or the presentation of credentials. Obligations are constraints on the future usage of

data, including the restriction of usages, e.g., playing a song at a quality below 80% if it has not been paid for, and action requirements, e.g., notification of a data subject whenever its data is accessed. Finally, *mechanisms* can be inhibiting (do not play a song at any quality if it was not paid for), modifying (reduce the quality to 80%), executing (issue a payment), and delaying (wait to see if the payment is not about to be received).

In this setting, what are the parameters to be negotiated? We do not think that the *data object* itself will be subject to (automated) negotiation. The consumer may directly ask for a concrete object or intensionally describe the object he is interested in (“an eBook on two-day hiking trips in the French Alps”). The seller will then likely set up a list of candidate objects together with their prices and quality attributes. The potential customer weighs content, quality and price against each other and picks one (or does not), according to his preference function. Where would back-and-forth negotiations enter the game? Since this question suggests that there is essentially no multiple-step negotiation when digital goods are traded, this is likely to provoke some disagreement (for instance, in the context of eCommerce, the literature suggests multiple rounds of negotiation in the case of auctions, for instance).

In terms of *policies*, consider an electronic library.² Books can be borrowed at a price of 5 Euros and must be deleted after 30 days, or at a price of 20 Euros and need not be deleted but must not be copied more than twice, or at a price of 30 Euros with unlimited rights. This again seems like a situation of live-or-let-die. However, once the customer has gotten these offers, he may want to get the right of three copies for 22 Euros. This seems unlikely to happen in automated settings. As a second example, consider waving certain privacy rights in exchange for a discount: the consumer gets 10% off the price if his name and address may be sold. It is not clear where negotiation would take place here either—it is a choice between two alternatives.

Finally, in terms of *enforcement mechanisms*, there indeed seems to be room for negotiation. Assume that a consumer can express the capabilities of his mechanisms. He could then send a list of his capabilities to the seller who would choose one and issue a respective licence (a policy becomes a licence when it is bound to a specific mechanism). However, the customer may not want to provide information on all his mechanisms but rather suggest one (the weakest), then a second, etc. Is this scenario realistic?

In sum, the main problem of automated negotiation in the context of access and usage control seems to be the definition of preference functions that allow a party to choose among alternatives. Does negotiation really take place?

²more precisely, the obligations part of policies: provisional actions such as payments have been discussed before, and access control requirements appear unlikely to be negotiated.