CS294-1 Algorithmic Aspects of Game Theory **Spring 2001**

Lecture 2: January 23

Lecturer: Christos Papadimitriou *Scribes: Andrea Frome(afrome@cs), Kunal Talwar(kunal@cs)*

2.1 Overview

We begin by looking at a set of theorems from various disciplines and how they relate to one another. From combinatorics, we take *Sperner's Lemma* which we can use to prove *Brouwer's Fixed Point Theorem* from topology. Brouwer's Fixed Point Theorem can be used to prove the *Arrow-Debreu Theorem* from economics which states that general equilibria exist, and can also be used to prove *Kakutani's Fixed Point Theorem*. Kakutani's Fixed Point Theorem can be used to prove that *Nash Equilibria* exist for all games. A graph illustrating how these theorems can be used to prove each other is given in Figure 2.1.

2.2 Sperner's Lemma

Here we consider an example application of Sperner's Lemma to a simplex in two dimensions, though the lemma can be generalized to higher dimensions. Take the following steps to set up the problem (see Figure 2.2):

1. Triangulate the simplex so that it's divided into smaller triangles.
2. Number the corners of the simplex 0, 1, and 2.
3. Label each of the vertices in the triangulation with either 0, 1, or 2, subject to the following rules:
    - a vertex on a side of the simplex cannot be assigned the same number as the corner of the simplex opposite that side, and
    - vertices inside the simplex can be labelled with any of the numbers.

Lemma 2.1. Sperner's Lemma: *If you label the vertices as described above, you will always have a small triangle somewhere in the simplex with its vertices labelled* $(0, 1, 2)$*. Moreover, you will have an* odd *number of such triangles.*

Proof. To prove Sperner's Lemma, we trace a path through the simplex that is guaranteed to end in a $(0, 1, 2)$ triangle. In order to make the proof cleaner, we first add some "triangles" to the $0-1$ side of the simplex (see Figure 2.3). We begin outside the simplex and enter the simplex using the boundary $0-1$ edge. Thus we enter a triangle that has at least one $0-1$ edge. At any point in the trace, if we are not in a $(0, 1, 2)$ triangle, the triangle has exactly two $0-1$ edges. We continue our trace by crossing the other $0-1$ edge. We cannot ever re-enter any triangle, since a triangle has at most two $0-1$ edges. Also, we cannot exit the

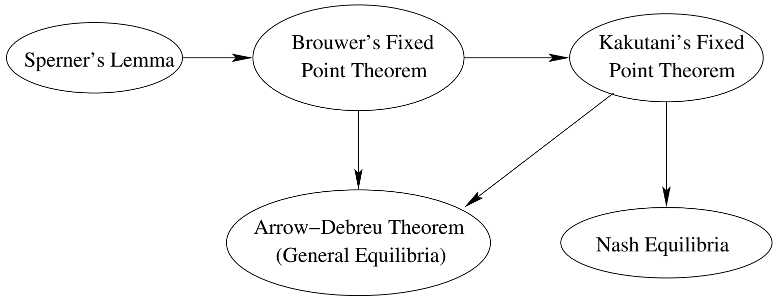

Figure 2.1: Graph showing which theorems can be used to prove others

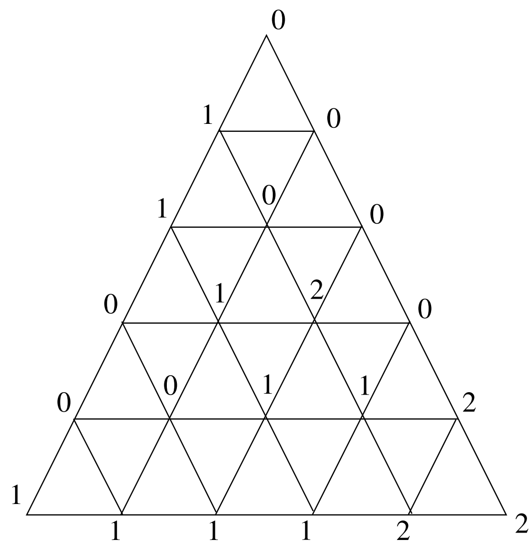

Figure 2.2: Example triangulated and labelled simplex

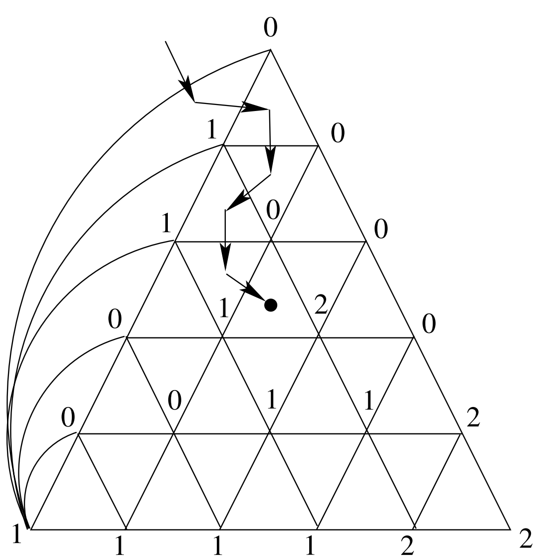

Figure 2.3: Example solution

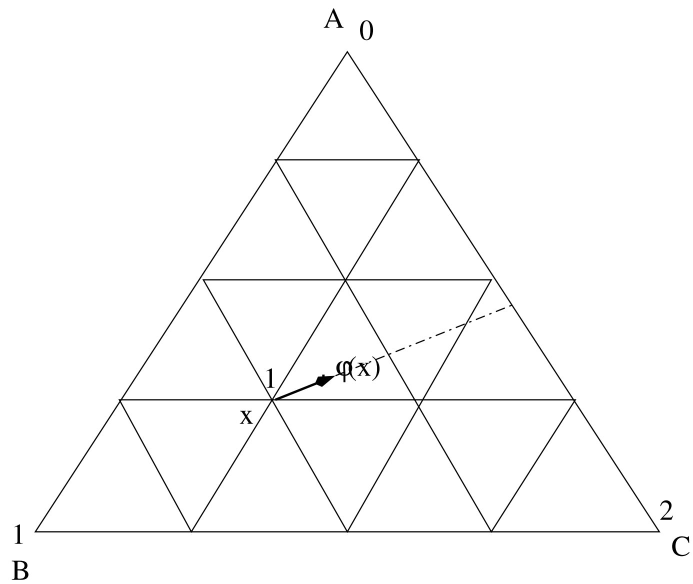

Figure 2.4: Labelling scheme

simplex because there is only one boundary $0 - 1$ edge. Since the triangulation has finitely many triangles, our trace must end at a traingle with exactly one $0 - 1$ edge, which has to be a $(0, 1, 2)$ triangle[1] □

Using Sperner's Lemma, we can now prove Brouwer's Fixed Point Theorem.

2.3 Brouwer's Fixed Point Theorem

Take an equilateral triangle—a two-dimensional simplex—and perform some transformation such that all points on the modified triangle lie somewhere within the boundary of the original triangle. Some possible transformations are rotating the triangle by some multiple of $\frac{\pi}{3}$, flipping the triangle along a bisector, or shrinking the triangle. (This does not include translating the triangle because some of the translated points would not lie within the boundary of the original.) After any such transformation ϕ, Brouwer's theorem says that there is a point x that is in the same position as it was in the original simplex.

Theorem 2.2. Brouwer's fixed point theorem (1915) *:Let S be any n-dimensional simplex and let $\phi : S \to S$ be any continuous function. Then ϕ has a fixed point, i.e. $\exists x^* \in S$ such that $\phi(x^*) = x^*$.*

Proof. We prove the theorem here for the special case where the simplex S is an equilateral triangle in $\mathbb{R}^2$. An essentially similar argument applies to the general case. The proof is a direct application of Sperner's lemma.

Let $T_0, T_1, \ldots, T_n, \ldots$ be a sequence of succesively finer triangulations (for example $Diameter(T_n) < 2^{-n}$). We define labellings $L_0, L_1, \ldots, L_n, \ldots$ for the triangulations as follows.For a vertex x in the triangulation, consider the vector from x to $\phi(x)$. Extend this vector until it meets a boundary of the simplex. If it crosses

[1] Consider each triangle in the triangulation as a room in a castle, and a $0 - 1$ edge as a door. Note that the castle has only one entry/exit, and each room has at most two doors. What the argument above says is that if we enter the castle, and move through the rooms, our path should end in a room with exactly one door.

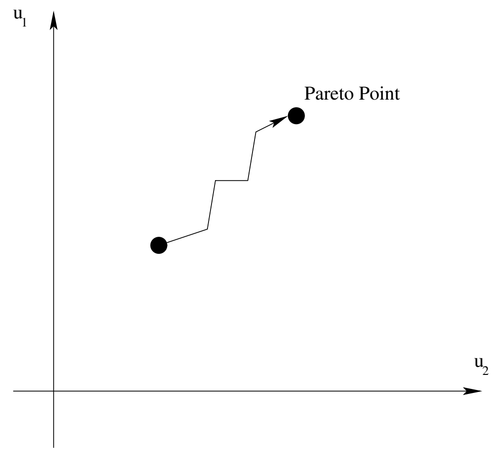

Figure 2.5: Graph of two people's utilities and one possible Pareto Point

edge AB (resp. BC, CA), we label X as 2 (resp. 0, 1). We break ties arbitrarily. It is easy to see that under this labelling, for any $\phi : S \to S$

- the vertices $0, 1, 2$ of the triangle get labels $0, 1, 2$ respectively.
- no vertex on the edge AB (resp. BC and CA), gets a label 2 (resp. 0 and 1).

Thus L_i is a legal Sperner labelling for trianglulation T_i. Thus by Sperner's lemma, there is a $(0, 1, 2)$ triangle, say t_i in T_i. Let m_i be the centroid of triangle t_i. Consider the sequence of points $\langle m_i \rangle$. This is a infinite sequence in S, and thus has a subsequence $\langle x_i \rangle$ that converges to a point $x \in S$ [2]. Thus x is arbitrary close to the centroid of some $(0, 1, 2)$ triangle.

We claim that x is a fixed point of ϕ. This is so because if $\phi(x)$ is different from x, we can find a small $(0, 1, 2)$ triangle containing x but not $\phi(x)$. Since ϕ is continuous, this triangle cannot be labelled $(0, 1, 2)$. For a more rigorous proof, the reader is referred to Appendix A. □

2.4 Arrow-Debreu Theorem (General Equilibria Exist)

To frame the problem, consider a marketplace with n agents a k different commodities. Each agent i comes to the market place with an endowment which we represent as a vector of goods or services: $\mathbf{e_i} \in \mathbb{R}_+^{\mathbf{k}}$. Every agent i has a utility function $u_i : \mathbb{R}_+^k \to \mathbb{R}_+$ which describes how much utility agent i gets from various amounts of the k goods. The utility function may be arbitrary, and it is not necessarily linear. Also we assume that the goods are infinitely divisible.

When agents come to the market, they want to exchange goods with other agents such that they improve their overall utility. There is a point in this exchange known as the *Pareto Point* where neither agent can improve their utility through further exchanges (see Figure 2.5).

The question is: How should exchanges be carried out to maximize the agents' utilities and at what prices? It's possible that the necessary sequence of exchanges is complicated or creates a cycle, and searching for an optimal exchange sequence is computationally intractible. Instead we could have everyone announce their

[2] This is a consequence of the Bolzano Wierstrass Theorem, e.g see Royden [RA98], page 153.

utility functions and try to solve for an optimum, but then we have no way of knowing whether people are being honest about their utilities.

A solution is to use set prices for each commodity $p_j \in \mathbb{R}^+$, which for now we will assume are "God-given". With prices in hand, there is no longer a dilemma; each agent buys a vector of goods $\mathbf{X_i}$ that maximizes $u_i(\mathbf{X_i})$ subject to

$$\mathbf{p} \cdot \mathbf{X_i} \leq \mathbf{p} \cdot \mathbf{e_i} \tag{2.1}$$

where $\mathbf{p} \cdot \mathbf{X_i}$ is the amount the agent pays for the goods she wants and $\mathbf{p} \cdot \mathbf{e_i}$ is the amount the agent received for her original endowment. This can be solved for $\hat{\mathbf{X}}_\mathbf{i}(\mathbf{p})$, the optimal purchase of goods for agent i when prices are $\mathbf{p}$. (We assume that $\hat{\mathbf{X}}_\mathbf{i}(\mathbf{p})$ is continuous and has a unique optimum.)

At some point the market may *clear*, meaning that all the goods for sale have been purchased. The total amounts of each commodity demanded by all agents is described by $\hat{\mathbf{X}}(\mathbf{p}) = \sum_{\mathbf{i=1}}^{\mathbf{n}} \hat{\mathbf{X}}_\mathbf{i}(\mathbf{p})$, and the market clears if

$$\hat{\mathbf{X}}(\mathbf{p}) \leq \mathbf{E} \tag{2.2}$$

where $\mathbf{E} = \sum_{\mathbf{i=1}}^{\mathbf{n}} \mathbf{e_i}$ is the total amount of goods brought to the market.

Theorem 2.3. Arrow-Debreu Theorem (1954): *There is always a price* $\mathbf{p}^*$ *such that* $\hat{\mathbf{X}}(\mathbf{p}^*) \leq \mathbf{E}$*, i.e. one can always find a price that clears the market. Such a price is a general equilibrium. This price is also a Pareto Point.*

(Note: While the price equilibrium is unique, there may be several Pareto Points, depending upon the agents' initial endowments.)

Surprisingly, finding $\mathbf{p}^*$ can be done with remarkably little communication. Consider the simplex of all prices

$$\sum_j p_j = 1.$$

If the price is such that there is excess demand, then the price should be increased by the amount that demand exceeds the endowment (assuming no inflation, we can normalize all prices such that they add up to one):

$$\Phi(\mathbf{p}) = \left\langle \ldots, \frac{\mathbf{p_j} + \mathbf{max}(\mathbf{0}, (\hat{\mathbf{X}}_\mathbf{j} - \mathbf{E_j})}{\text{normalization constant}}, \ldots \right\rangle \tag{2.3}$$

Either the endowment is greater than or equal to what the agents want in which case the price stays the same, or the price is increased by the difference between what the agents want and the total endowment. This transformation maps a vector of prices to another point in the simplex, and Brouwer's Fixed Point Theorem says that there must always be a fixed point to such a tranformation. Call the fixed point $\mathbf{p}^*$.

Claim 2.4. *At fixed point* $\mathbf{p}^*$*, for every product* j*,* $\hat{X}_j \leq E_j$*.*

Proof. Let $\mathbf{d} = \hat{\mathbf{X}} - \mathbf{E}$. Let $d'_j = max(0, d_j)$. We want to show that $d'_j = 0$ for all commodities j.

We have from 2.3,

$$\Phi(\mathbf{p}) = \frac{1}{k}(\mathbf{p} + \mathbf{d'}) \tag{2.4}$$

for some normalization constant $k > 1$.

Moreover, adding equations 2.1 for all agents i, we get

$$\mathbf{p} \cdot \hat{\mathbf{X}} \leq \mathbf{p} \cdot \mathbf{E} \tag{2.5}$$

i.e.

$$\mathbf{p} \cdot \mathbf{d} \leq \mathbf{0} \tag{2.6}$$

Since $\mathbf{p}^*$ is a fixed point, $\Phi(\mathbf{p}^*) = \mathbf{p}^*$. Taking dot product of equation 2.4 with $\mathbf{d}$, we get,

$$\mathbf{p}^* \cdot \mathbf{d} = \frac{1}{k}(\mathbf{p}^* + \mathbf{d}') \cdot \mathbf{d}$$

Rearranging, and noting that $\mathbf{d'} \cdot \mathbf{d} = \mathbf{d'} \cdot \mathbf{d}$, we get

$$\|\mathbf{d}'\|^2 = \mathbf{d}' \cdot \mathbf{d}' = (k-1)\mathbf{p} \cdot \mathbf{d} \leq \mathbf{0}$$

Thus $\mathbf{d'}$ must be zero. □

2.5 Nash Equilibria Exist

The proof of Nash's theorem requires the use of another fixed point theorem more general than the Brouwer's fixed point theorem.

Theorem 2.5. Kakutani's fixed point theorem *Let $\Phi : S \to 2^S$ be any convex valued function such that for any sequence $\langle (x_i, y_i) \rangle$ converging to (x, y), if $\forall i \in \mathbb{N} : y_i \in \Phi(x_i)$, then $y \in \Phi(x)$ (this property is called* graph continuity*). Then Φ has a fixed point, i.e. $\exists x^* \in S$ such that $x^* \in \Phi(x^*)$.*

Proof. We give here a proof sketch for the 2-d simplex case. As in the proof of Brouwer's theorem, we consider a sequence of succesively finer triangulations $T_0, T_1, \ldots, T_n, \ldots$. For a triangulation T_i, we define a function ϕ_i as follows. For a vertex x of the triangulation, $\phi_i(x)$ is set to some $y \in \Phi(x)$, chosen arbitrarily. For any other point $x \in S$, $\phi_i(x)$ is defined by linear interpolation in the triangle containing x. Thus ϕ_i is a continuous map. By Brouwer's fixed point theorem, ϕ_i has a fixed point x_i^*. Now consider the sequence $\langle x_i^* \rangle$. It has a convergent subsequence, that converges to some point $x^* \in S$. Using the graph continuity of Φ, it can be shown that $x^* \in \Phi(x^*)$. □

For a more rigorous proof of Kakutani's fixed point theorem, the reader is referred to [KA41].

For any game, a Nash equilibrium suggests that none of the players has an advantage in changing his strategy, without the other players changing their strategy as well. More formally,

Definition 2.6. *Let S_1 and S_2 be the set of possible strategies for player 1 and 2 respectievely. Let $p_1 : S_1 \times S_2 \to \mathbb{R}$ (resp. p_2) be the payoff function for player 1 (resp. player 2). Then a tuple $(s1, s2)$ is said to be in Nash equilibrium if*

$$\begin{aligned} \forall s_1' \in S_1 : \quad & p_1(s_1, s_2) \geq p_1(s_1', s_2) \\ \forall s_2' \in S_2 : \quad & p_2(s_1, s_2) \geq p_2(s_1, s_2') \end{aligned}$$

While the above definition captures all pure strategies, it does not consider the possibility of *mixed strategies*, i.e. strategies where a player chooses a move $s \in S$ randomly according to some distribution Π on S. When we allow mixed strategies, we assume that each player tries to maximize his expected payoff. Thus if player i chooses from a distribution Π_i on S_i, each player tries to maximize

$$p_i(\Pi_1, \Pi_2) = \sum_{s_1 \in S_1, s_2 \in S_2} \Pi_1(s_1)\Pi_2(s_2)p_i(s_1, s_2)$$

A tuple of mixed strategies is said to be in mixed Nash equilibrium if

$$\begin{array}{llcl} \forall \text{ prob.distributions } \Pi_1' \text{ on } S_1: & p_1(\Pi_1, \Pi_2) & \geq & p_1(\Pi_1', \Pi_2) \\ \forall \text{ prob.distributions } \Pi_2' \text{ on } S_2: & p_2(\Pi_1, \Pi_2) & \geq & p_2(\Pi_1, \Pi_2') \end{array}$$

While for many games, it is the case that no pure Nash equilibria exist, mixed Nash equilibria always exist.

Theorem 2.7. Nash's theorem: *Every game has a mixed strategy Nash equilibrium.*

To prove Nash's theorem, let $\Phi_1(\Pi_2)$ be the set of all mixed strategies $\Pi_1' \in \mathcal{D}(S_1)$ that maximize $p_1(\Pi_1', \Pi_2)$ for a mixed strategy $\Pi_2 \in \mathcal{D}(S_2)$ of player 2. Similarly define the set $\Phi_2(\Pi_1)$. Now let

$$\Phi(\Pi_1, \Pi_2) = \Phi_1(\Pi_2) \times \Phi_2(\Pi_1)$$

It is easy to see that $\Phi : \mathcal{D}(S_1) \times \mathcal{D}(S_2) \rightarrow 2^{\mathcal{D}(S_1) \times \mathcal{D}(S_2)}$ as defined above is convex valued and graph continuous. Hence by Kakutani's fixed point theorem, Φ has a fixed point (Π_1^*, Π_2^*). This, by definition is a mixed Nash equilibrium.

2.6 Overview of Course Topics

A brief overview of the topics for the semester:

1. Background on equilibrium theorems.
2. Notions of Equilibria.

   Nash equilibria always exist, but often there are too many of them. Much of game theory has been a critique of Nash equilibria.
3. Games played by automata.

   To explain people's strategies in Prisoner's Dilemma, we need a more refined theory. Perhaps repetition and reputation are the the key. If we consider resource bounds, then automata become relevant to game theory; people don't have infinite reasoning capacity and we can model them using automata with limited states.
4. Evolutionary Game Theory

   Discuss John Maynard Smith. Talk about game theory in the context of evolution, finding an evolutionarily stable strategy, e.g. "tit for tat".
5. Mechanism Design

   Inverse game theory. Given desired outcomes, design a game so that agents acting rationally will behave in a desired way. An example is the Vickero auction where the highest bidder is the winner, but the winner pays the amount of the second highest bid.

6. Fairness.
7. Auctions, including combinatorial auctions.
8. Price of Anarchy

   Take a network where it is possible to control what route each packet takes to its destination. If a central authority were to determine an optimal routing policy given all the information in the network, the improvement in performance would be at most twice over a system where each agent optimizes its own packets.

References

[1] H. L. Royden, *Real Analysis*, 3^{rd} edition, 1998, Prentice Hall.

[2] S. Kakutani, "A Generalization of Brouwer's Fixed Point Theorem", *Duke Mathematical Journal* 8, 1941, pp.457-459.

Appendix

A Details in proof of Brouwer's theorem

Theorem 2.8. *Consider the labelling defined in theroem 2.2. Let x be the limit of a sequence $\langle x_i \rangle$ of centroid of a sequence of $(0, 1, 2)$ triangles with diameter converging to 0. Then x is a fixed point of ϕ.*

Proof. Assume the contrary, i.e. $\phi(x) \neq x$. Then $t = \|\phi(x) - x\| > 0$. Let $\epsilon = t/2^k$ for some k. Since ϕ is continuous and $\epsilon > 0$, there exists $\delta > 0$ such that $\forall y \in N(x, \delta) : \phi(y) \in N(\phi(x), \epsilon)$.

Let $\epsilon' = \frac{1}{2}\min(\epsilon, \delta)$. Since x_i converges to x, $\exists N_1 \in \mathbb{N}$ such that $\forall n > N_1 : x_n \in N(x, \epsilon')$. Further, since the diameter of the triangulations converges to zero, $\exists N_2 \in \mathbb{N}$ such that $\forall n > N_2 : diameter(T_n) < \epsilon'$. Thus for $n > \max(N_1, N_2)$, a $(0, 1, 2)$ triangle t_n lies wholly inside $N(x, 2\epsilon') \subseteq N(x, \delta)$. Thus by the continuity condition above, each of the three vertices a, b, c of the $(0, 1, 2)$ triangle maps to some point in $N(\phi(x), \epsilon)$, outside the triangle. Since the triangle is a $(0, 1, 2)$ triangle, the vectors $\phi(a) - a, \phi(b) - b, \phi(c) - c$ span an angle of at least $\pi/3$ (by our labelling scheme). However, for k large enough, the angle spanned can be made arbitrarily small (see figure 2.6 (b)), which is a contradiction. Hence the claim. □

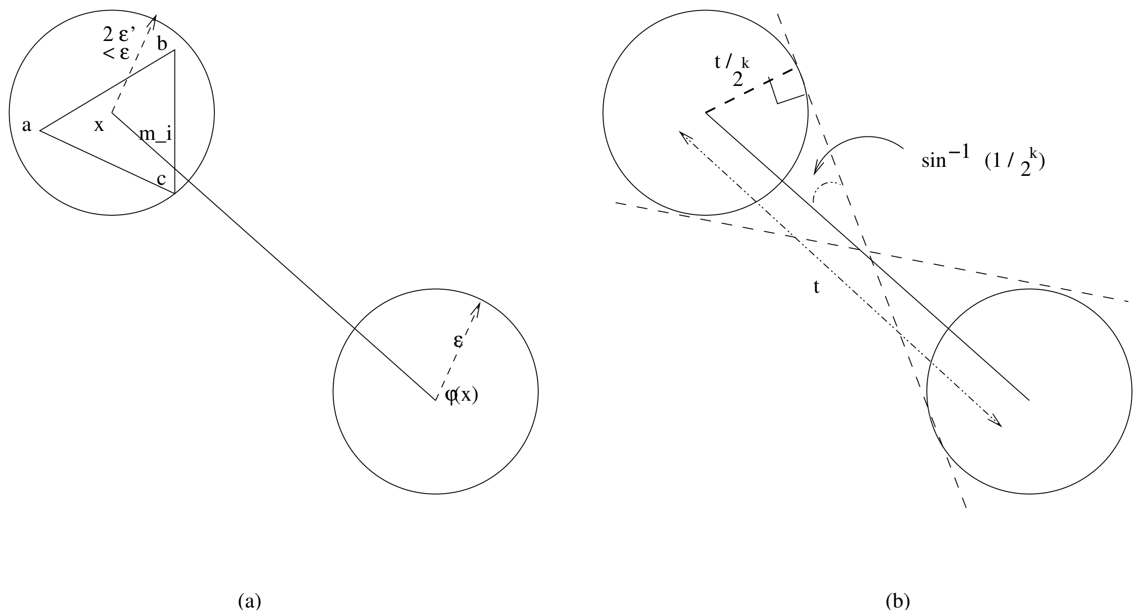

Figure 2.6: (a) Proof details (b) Angle spanned by vectors should be small